

Report on the VLDB Workshop on Technologies for E-Services (TES)

Cairo, Egypt
September 14 and 15, 2000

Fabio Casati, Umesh Dayal, and Ming-Chien Shan
Hewlett-Packard Labs
E-Mail: [casati,dayal,shan]@hpl.hp.com

1 Aims and scope

The Web is changing every aspect of our lives, but no area is undergoing as rapid and significant changes as the way businesses operate. Today, large and small companies are using the Web to communicate with their partners, to connect with their back-end systems, and to perform electronic commerce transactions. The next chapter of the Internet story is the evolution of today's e-business and e-commerce systems into "e-services", such as order procurement, on-line trading, customer relationship management, product promotion, or real-time car navigation and traffic information services. In order to make e-services available to customers, service providers need to address several issues, such as:

- **e-service description:** which are the attributes of an e-service that should be made visible to customers or applications, and how they should be described.
- **e-service advertisement:** how service providers can publish service description so that they can be discovered and accessed by customers and applications.
- **e-service discovery and selection:** how customers and applications can discover and select the e-service (or the combination of e-services) that best fulfill their requirements.
- **e-service composition:** how basic e-services (possibly offered by different companies) can be combined to form value-added,

reliable services. Which architectures, models, and languages can achieve zero latency service integration and cross-organizational business process automation.

- **e-service delivery:** how e-services are delivered to businesses and to customers
- **e-service monitoring and analysis:** how service executions can be monitored and how service execution data can be analyzed in order to improve the service quality or efficiency
- **e-service contracts:** how to agree on and perform legal contracts between service providers and clients electronically.
- **e-service ratings:** how to validate service claims and evaluate the quality of the different service providers.

The goal of this workshop has been to identify the technical issues and the infrastructures that enable organizations to provide e-services to businesses and customers, either directly or through market exchanges or eBusiness portals. The workshop brought together more than fifty researcher and developers working in the areas of e-commerce, Web portals, application servers, ontology, description languages, and process management systems.

The workshop program included four industry talks and sixteen papers organized into the four areas of customization and interaction models, service composition, service

infrastructure and methodologies, and web application data management.

In the following, we will briefly illustrate the content of the accepted papers in each of these areas, and we will next summarize the four industry talks.

2 Customization and interaction models

Mass customization is a fundamental feature for every e-business solution. Several papers in the workshop proposed models and technologies for supporting personalized content and service delivery. These papers were grouped in the first paper session.

In *Generation and Verification of Heterogeneous Purchase Processes*, by Popovici, Schuldt, and Schek of ETH Zurich, the authors show how to model complex interactions among customers and suppliers in performing e-commerce transactions. In particular, the paper focuses on the problem of managing and merging heterogeneous purchase protocols, and present the *purchase coordinator*, a prototype system that manages such complex, heterogeneous transactions.

The paper *Integrating and Customizing Heterogeneous E-commerce Applications* presents an XML- and UML-enabled system, called *Application Manifold*, that aims at simplifying the integration and customization of e-commerce applications. The system is an application generator that operates according to specifications given in a declarative language. The system deals with all the aspects of the integration, including services, roles, application flow, and data.

Web Services and Information Delivery for Diverse Environments, by Freire and Kumar of Bell Labs, discusses the issues involved in making existing Web content and services available for diverse environments, and describe *PersonalClipper*, a system that allows users to create customized (and

simplified) views of Web sites that are suited for different types of terminals.

The paper *Toward a Comprehensive Model of the Content, Structure, and User Interaction of a Web Site*, by Datta, VanderMeer, Navathe, (Georgia Tech) and Ramamritham (Umass) tackles the problem of delivering personalized e-services which are sensitive to the current and past user behaviors. In particular, the paper proposes a model of a web site that includes not only the site content and organization, but also user interaction with the site.

Finally, in *Active Rules for XML: A New Paradigm for E-Services*, by Bonifati, Ceri, and Paraboschi (Politecnico di Milano), the problem of active document management is discussed. The paper focuses on the definition of active rules for document management and specifically on the design of an event model for XML documents and its integration with XML servers. The authors show how events can be specified in the context of XSL and of Lorel (a query language for XML documents). The authors also demonstrate, through simple examples, that active rules can be very effective for the implementation of e-services, and outline the main issues that will be faced in order to support active document management on XML servers.

3 Service composition

The area of service composition is attracting strong interests from both the academia and the industry. This is demonstrated by the high number of related submissions and by the focus of three industry talks. Four papers were finally selected in the service composition space.

XML-enabled Workflow Management for E-Services across Heterogeneous Platforms (by Shegalov, Gillmann, and Weikum of University of the Saarland, Germany) proposes an XML-enabled

architecture for distributed workflow management, implemented on top of the *Mentor-lite* prototype. A crucial component of the architecture is the *XML mediator*, that handles the exchange of business and flow control data between workflow and business-object servers on one side and client activities on the other side via XML messages over http. The major benefit of the architecture is that it integrates e-services with scalable efficiency and very little explicit coding.

The paper *Process-based E-Service Integration*, by Georgakopoulos, Cichocki, and Baker (Telecordia), and Schuster (Heyde AG), presents a model for defining multi-enterprise supply chains. The model supports *service activities* for modeling the services themselves, primitives for composing services into supply chains, and primitives for automating service coordination. One of the most important aspects of the proposed approach is the decoupling between service interface and service implementation. This enables processes to include *placeholders*, i.e., activities specified only by an abstract interface. Placeholders are substituted during runtime with one of available services that implement the specified interface. Service selection is done through a semantic broker that has knowledge about the service quality and capabilities.

E-ADOME: a Framework for Enacting E-Services, by Chiu (Dickson Computer Systems), Karlapalem (Hong Kong Univ. Of Science and Technology), and Li (Univ. of Hong Kong) proposes a service composition framework that includes automated resolution of expected exceptions, cooperative exception handling, and user-driven computer supported resolution of unexpected exceptions. The framework is supported by and implemented on top of the workflow system E-Adome, so to leverage its functionality in modeling and managing exceptional situations.

The paper *Composing and Maintaining Web-based Virtual Enterprises*, by Benatallah (Univ. of New South Wales, Australia), Medjahed and Bouguettaya (Virginia Tech), Elmagarmid (Purdue Univ.), and Beard (CiTR Ltd, Australia) presents a generic framework for creating and maintaining virtual enterprises (VEs). The authors introduce a language, called *WebBIS-SDL*, to cater for the definition and maintenance of VEs. To show the feasibility of the proposed framework, the authors have implemented a CORBA-based prototype that allows easy definition and maintenance of VEs.

4 Methodologies and infrastructures

The third paper session of the workshop was centered around methodologies for *developing* e-services and infrastructure for *delivering* e-services. The session was opened by the paper *Designing Components for E-Services*, by Mecella (Univ. La Sapienza, Roma, Italy) and Pernici (Politecnico di Milano, Italy). The authors discuss a component-based approach for the development of distributed applications. In particular, the paper focuses on the notion of component in the framework of e-service design, especially when these services are based on legacy systems. The authors give definition of *stateful* vs *stateless* components and discuss service characteristics and applicability in the different stages of the application development.

In *Enterprise Modeling and Integration in the M-Complex Approach*, by Berio and Di Leva (Univ. of Torino, Italy) a methodology for complex system development is presented. The authors introduce the M*-COMPLEX approach to study complex systems, based on authors' previous work about the M*-OBJECT methodology for the design of highly complex information systems. M*-COMPLEX is partially supported by a software tool, called M*-PROCESS,

which allows the representation of the business knowledge at multiple levels of abstraction and from multiple points of view. The main methodological phases are presented through a case study related to VE modeling.

Finally, the paper *Matching and Deriving Dynamic Constraints for E-Commerce Negotiations*, by Mizuho Iwaihara (Kyushu Univ., Fukuoka, Japan), focuses on the issue of negotiation, dealing in particular with cases that involve the exchange of complex conditions. The author presents *derived dynamic constraints*, which define complex constraints using *dynamic constraint algebra* queries. The authors provide several examples showing matching of buyer and supplier constraints, even when incomplete specifications are involved.

5 Web application data management

The final session of the workshop was centered on data management for e-service applications. In *Smart Supply Web: An Application of Web-based Data and Workflow Mediation*, by Hull, Kumar, Simeon (Bell Labs, New Jersey), the authors present their vision of a *Supply Web* as the next major advance in the e-commerce world in terms of how e-services are delivered. After highlighting shortcomings of existing supply chain solutions, the paper proposes an architecture to show how three important technologies, namely, workflow systems, data integration and on-line decision support, will have to come together to enable seamless deployment of e-service supply web solutions. After discussing research efforts that support the three technologies separately, the authors describe the challenges of integrating them effectively.

ObjectGlobe: Ubiquitous Query Processing on the Internet (by Braumandl, Keidl, Kemper et al., of Passau Univ., Germany) presents a

distributed and open query processor. The goal of the *ObjectGlobe* project is to establish an open marketplace in which *data* and *query processing capabilities* can be distributed and used by any kind of Internet application. The goal is to make it possible to execute virtually any kind of query operator on any machine and any kind of data on the Internet. One of the main challenges in the design of such an open system is to ensure privacy and security. The authors also introduce the *ObjectGlobe* security requirements and show how basic components such as the optimizer and runtime system need to be extended in order to handle them.

The session was closed by the paper *The Cameleon Web Wrapper Engine*, by Firat, Madnick, and Siegel (MIT Sloan School of Management, Cambridge, MA). In this paper, the authors describe a technology and infrastructure to support the flow of information among services on the web and their interconnection with legacy systems that were designed to operate with traditional relational databases. The proposal is centered around a web wrapper engine called *Caméléon*. *Caméléon* extracts data from web pages using declarative specification files that define extraction rules. *Caméléon* is based on the relational model and is designed to work as a relational front-end to web sources. Query results by *Caméléon* are presented in either XML or HTML. Users can also easily call *Caméléon* from other applications.

6 Industrial program

The workshop included four industry talks. Dick Tsur, of BEA Systems, gave an overview of the e-commerce solutions proposed by BEA. In addition, he outlined the recently established *Office of Technology* within BEA. Its purpose is to assess new technologies of potential significance to the company and to foster external relationships with potential partners in technology development.

The second talk was by Christoph Bussler (Oracle) on *The Role of WFMSs in EAI, B2B and ASP Environments*. The presentation stressed the crucial role that Workflow Management Systems (WfMSs) play in driving integration and interaction within and across corporations. Christoph also introduced several examples showing the use of WfMSs in EAI, B2B and ASP environments and discussed the main functionality provided for advanced systems integration.

Johannes Klein (Microsoft) presented the *BizTalk Orchestration Architecture*. The presentation focused on an architecture for Internet-based business processes spanning across organizational boundaries. The concepts have been illustrated through an example focused on the corporate procurement application.

The final presentation, by Sekhar Sarukkai and Fabio Casati (Hewlett-Packard) described the service composition features provided by *e-speak*. *e-speak* is a platform for developing, deploying, and delivering e-services. The talk presented a particular (meta-) service that enables the definition of complex services by composing existing ones. Once defined, the composite services are registered on the *e-speak* platform and made available to customers just like any other service.

7 Conclusion

The number (and quality) of submissions and the large participation (which exceeded the organizers' expectations) has shown that both industry and academia have a great interest on e-services.

Most of the submissions focused on e-service *personalization, composition, development methodologies, and application and data management*. While some products and solutions exist in these spaces, the general consensus has been that more needs to be done in these areas in order to effectively address the problems.

On the contrary, very few submissions were related to service description languages and interaction protocols. This is probably an indication that the community believes that existing approaches and standardization efforts are going in the right directions, and that the main problem is actually to avoid the proliferation of the different description languages and communication protocols.

Acknowledgement

The TES 2000 organization would like to thank those who helped in arranging the workshop as well as all reviewers, authors, presenters, and participants for enabling and contributing to the interesting and lively discussions. We would like in particular to help our co-organizers Alex Borgida, Christoph Bussler, Stefano Ceri, Dimitrios Georgakopoulos, and Tamer Ozsu.