



# Models and Languages for Describing and Discovering E-services

Fabio Casati and Ming-Chien Shan  
Hewlett-Packard Laboratories  
1501 Page Mill Road, 1U-4  
Palo Alto, CA, 94304 USA  
[casati,shan]@hpl.hp.com

## 1. INTRODUCTION

The next chapter of the Internet story is about the evolution of today's various software applications into *e-services*. An e-service is a business function made available via the Internet by a service provider, and accessible by clients that could be human users or software applications. Today, many e-services are already available on the Internet. Examples of e-services include delivery of customized news and stock market updates, on-line ticket reservations, flight status monitoring and changes, and e-advertisement. It is expected that the number and type of on-line services is likely to increase at a fast pace in the near future.

The main benefit of the e-services environment is that clients are able to dynamically discover the available e-service that best meets their needs, to examine its properties and capabilities, and to determine if and how to access it.

However, in order to deliver e-services to clients, service providers are faced with several challenges. In particular, they need to *describe* e-services in a way that is accessible and understandable by the clients. The description should include meta-information about the service as well as information about the specific functions offered by the service. In addition, e-service descriptions must specify how clients can interact with them. In addition, in order to achieve their business goals, clients will typically have to invoke several functions. For instance, to rent a car, clients may first *browse* the set of available cars, then *check* prices and availability, and finally *make a reservation*. Hence, the description of an e-service should formally define the allowed interactions and the execution dependencies among these functions, possibly according to some standard ontology.

E-services also need to be *advertised* in web directories, so that they can be *discovered* by brokers as well as by end-users. Service directories should be structured in such a way that they can always provide up-to-date information. This is very important, since the e-services world is highly dynamic.

In order to support e-service delivery, software vendors and standardization bodies are defining languages and protocols that facilitate service description and discovery. In this tutorial we discuss the requirements for such languages and protocols, and we present solutions proposed by individual vendors as well as by industry consortia. We start the tutorial by introducing the notion of e-service, and by presenting basic concepts and terminology. In particular, we will introduce the "classic" distinction between B2B, B2C, P2P, etc., and we will discuss direct e-service delivery versus delivery mediated by e-services hubs.

With respect to e-service description, we will show which characteristics of e-services need to be modeled so that clients can understand the service provided and learn how to interact with it. For instance, a service description could include the specification of the service attributes, of the message exchange format, or could include constraints over sequences of interactions, i.e., specify the *conversations* that can be carried on with the e-service. Next, we present different approaches for describing e-services. We introduce vendor-specific languages and models, as well as languages and models proposed by industry consortia and standardization bodies. In particular, we highlight similarities and differences between approaches, and discuss their applicability and key success factors.

Once a service has been described, it must be advertised so that it can be discovered. In the tutorial, we will discuss the challenges involved in advertising and discovering a service over the Internet. We first present the main issues and motivate why there is the need for solutions in this space. Then, we introduce the languages and technologies that can help solving this problem, and discuss applicability and alternatives. As for e-service description, we present both vendor-specific approaches as well as languages proposed by consortia or standardization bodies.

In particular, in this part of the tutorial we focus on UDDI, a standard supported by almost all the leading IT vendors. UDDI defines interfaces for e-service directories that can be accessed by humans or applications in order to search the service of interest. In addition, UDDI specifies the requirements that operators implementing UDDI directories should satisfy, especially concerning the global replication of data across all UDDI directories.

We conclude the tutorial by presenting our vision of what lies ahead in the e-services space and by discussing some open research issues.

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 and/or a fee.

ACM SIGMOD 2001 May 21-24, Santa Barbara, California, USA  
Copyright 2001 ACM 1-58113-332-4/01/05 ...\$5.00.