



Universal Description, Discovery and Integration

UDDI Executive White Paper

November 14, 2001

The Promise of Service Centric Computing

With the advent of service-centric computing, the Internet presents incredible value and reach for businesses of all sizes, providing opportunities to find new customers, streamline supply chains, provide new services, and secure unprecedented financial gain.

Organizations that have decisively moved their business online are already realizing significant economic and competitive gains: increased revenue, lowered costs, new customer relationships, innovative branding opportunities, and the creation of new lines of customer service.

Despite the promise of service-centric computing, a major impediment has held back its enormous potential. Most eCommerce-enabling applications and Web services currently in place take divergent paths to connecting buyers, suppliers, marketplaces, and service providers.

For example, without large investments in technology infrastructure, businesses, such as a semiconductor manufacturer in Taiwan or a cabinetmaker in Georgia can transact Internet-based business only with the global trading partners they have discovered and, of those, only the ones using the same applications and Web services.

In order to fully open the doors to these players, truly successful service-centric computing requires that businesses be easily able to discover each other, to make their needs and capabilities known, and to integrate services using each businesses' preferred technology, Web services, and commerce processes.

Challenges

Until now, there has been no central way to easily get information about what standards different services support and no single point of access to all markets of opportunity, allowing them to easily connect with all potential service consumers.

For instance, successful B2B eCommerce requires seamless access to information about trading partners and the ability to integrate with them. However, with the myriad ways to describe products and Web services – and because global eCommerce participants have not yet agreed on one standard or backbone on which to communicate their services – finding and working with potential trading partners is severely limited. This fundamental challenge is limiting the promise of business to business collaboration on the web, making it harder for buyers to get return on their eCommerce investment and for all B2B participants to easily add trading partners and services.

This basic problem also applies to the internal IT issues often faced by organizations. Often large, complex IT infrastructures will have a number of services that do not know of each other, let alone recognize how to integrate to leverage complimentary assets.

Solutions

Marketplaces, businesses, and directory providers are all attempting to solve these communication and transaction problems, and all are adopting distinct and divergent approaches centered on their own requirements. The result is a staggering diversity in approach, content, and architecture that is preventing the optimum utilization of service-centric computing by businesses of all sizes around the world.

Web services technology based on standard methods of description, discovery, and integration alleviate the implications of this diversity. Web services are self-contained, modular business applications that have open, Internet-oriented, standards-based interfaces. Web services communicate directly with other Web services via standards-based technologies. These standards-based communications allows Web services to be accessed by customers, suppliers, and partners independent of hardware, operating system, or even programming environment. The result is an

environment where businesses can expose their current and future business applications as Web services, which can be easily discovered and consumed by external partners. Web services offer improved time-to-integration and lower overall cost-of-ownership as compared to today's EDI and B2B solutions.

Web services enable businesses to:

- Improve collaboration with customers, partners, and suppliers by reducing integration time and expense compared to existing EAI and B2B solutions. Results in reduced inventory and transaction costs and improved supply chain efficiency.
- increase revenue via expanded distribution channels, quicker time-to-market for new value-added services, and enabling public discovery of existing assets
- enhance customer service levels by allowing customers and trading partner access to core systems
- generate new revenue opportunities through creation of private trading networks, or eco-systems
- quickly respond to changing market conditions and customer preferences by utilizing loosely-coupled, modular services
- protect investments and avoid technology lock-in via a standards-based architecture and hardware, operating system, and programming language neutrality

The last few years have seen remarkable evolution in Web-based B2B eCommerce, electronic sales, online auctions, dynamic electronic marketplaces, and applications that process and route information. These comprise the essential foundation of B2B eCommerce infrastructure, ensuring an organization's ability to establish connectivity, put product or service information online, access and interact with a broad range of customers, process transactions, and fill orders.

To accomplish this, a comprehensive solution is needed for businesses to publish their information to any customer or business partner around the world. Just as a common method for publishing data on the web spawned the evolution of e-business, a common means to publish information about business services will make it possible for organizations to quickly discover the right trading partners out of the millions that are online; to define how to conduct business once preferred businesses are discovered; and to create an industry-wide approach for businesses to quickly and easily integrate with their customers and partners on the Internet with information about their products and services, and how they prefer to be integrated into each other's systems and business processes.

New Approach

In order to understand the functionality provided by UDDI, it is important to review the type of Internet registry that most of us have been using for the last 5 years or so -- the Internet search engine. Internet search engines are large databases/registries that you use as an index into web sites on the Internet. Technically, they are just giant databases of URL's and keywords associated with those URL's. The primary focus of Internet search engines is to allow a human being sitting at a web browser to search for and find web sites containing the content they are looking for. Search engines are hard-wired to serve up only URL's -- and those URL's are expected to be used by web browsers.

So how do Search engines get populated with URL's and keywords for finding them? Well, there are a few ways this happens. Search engines can employ "web crawlers" that go out and find web sites, make note of their URL's, and store any keywords and other data that could be used to index the URL's of the sites. If you have a web site and you hope it can be found using an Internet search engine, you might -- if you are lucky -- have it be found by one of these web crawlers. The problem is you don't know if the web crawlers find all the sites, you don't know how frequently they are run, etc. It's a pretty inexact way to get your URL published into a search engine.

Another way you can get your URL published into a search engine is to strike up a business deal with the company running the search engine. For example, you might pay them to list your URL in their engine. And if you pay them a little more, you might be able to ensure that your URL is one of the first ones listed when a user browses for the type of content you are providing on your site. And there are

other types of business arrangements that could ensure you are listed on a search engine. For example, if you advertise with a search engine or an affiliate company of theirs, you also could earn the right to have your URL listed.

So the way an Internet search engine is used by a client is really simple. If you are the client, you first enter the URL of the Search engine's web site into your browser. Then you type in the keyword or keywords that you think might be the way to find the content you want. At this point you are hoping that what you typed in was used by the search engine to index the type of web sites you want. You might be right: you might be wrong: you hope you are right.

After entering your search text, you click on a "Go" button or a "Search" button or whatever button they have for you to click to execute the search. Clicking the button submits your search criteria to the Internet search engine. This causes a search into its database for all URL's that have been associated with the text string you typed in. All the matches that are found are sent back to your browser as hyperlinks in a results page. Then you can click on these hyperlinks to discover if the web sites they point to are sort of what you were interested in.

An interesting point is that all you can get back from these search engines are URL's to web sites. You don't get back email addresses or FTP sites. Just URL's to web sites. And all those URL's serve up only one type of documents -- HTML documents. They don't serve up SOAP documents or any kind of XML documents. They are only intended to serve documents that can be read by web browsers.

Publishing your URL's to the UDDI registry is much more of an exact science than it is trying to get your web site's URL into an Internet search engine. With UDDI, you have complete control over which of your business and service and service address information is published and when. This is because you are the one who publishes it.

How do you publish your businesses and services to UDDI? There are a number of ways to do it. All UDDI Operators (such as HP, IBM, SAP, and Microsoft) are required to provide a web site that you can point your browser to and use to add your businesses, services, and service addresses. You can also use richer, thick-client applications that can run on your local computer and can access the UDDI registry over the Internet using SOAP/XML messages. And your programs can register themselves without any human intervention using SOAP/XML messages and one of the freely-available client-side Java, Visual Basic, C#, or COM API's for accessing UDDI registries.

The important difference between how you publish to UDDI and how you publish a web site's URL into an Internet search engine is that you do it yourself -- you have complete control. Another important distinction, is that your service addresses can be not only URL's to web sites. They can also be email addresses, FTP sites, and any other type of computing resource that can be accessed over the Internet and can be addressed using a URI.

A common analogy used for UDDI is a "phone book for Web services." It has business names, business mailing addresses, contact names, contact phone numbers, Web services offered by businesses, addresses of Web services, meta-data describing the "interfaces" of Web services, etc. Lots of information.

Much of the data stored in the UDDI registry is useful for fully automated computing end points that want to find other end points that are offering Web services and then run those services. But a lot of the information is also useful for humans. For example, software developers who want to leverage existing Web services that have been published to UDDI could use a UDDI Browser to peruse what is "out there" for their consumption. Since UDDI is like a giant phone book, other people may want to use it just to find the contact information for certain businesses. There are many ways that human beings will be able to leverage the contents of the UDDI registry.

Currently, the richest interfaces that people can use to interact with the data in the UDDI registry are the web sites provided by the registry Operators. However, the Web services marketplace is starting to see richer, client-side GUI-based tools. These tools can provide the added functionality that a full-fledged GUI application can bring while at the same time tying back to the immense wealth of data in the UDDI registry.

UDDI Project

The advent of Web services has carried deep implications for business processes and organizational culture. UDDI is a **major** advance - the first cross-industry effort driven by platform providers, software developers, marketplace operators, and business leaders that comprehensively addresses the problems limiting the growth of service-centric computing, and that will benefit businesses of *all* sizes by creating this global, platform-independent, open framework to (1) discover each other, (2) define how they interact over the Internet, and (3) share information in a global registry that will more rapidly accelerate the global adoption of service-centric computing.

UDDI is a building block to enable businesses to quickly, easily and dynamically find and transact with one another via their preferred applications. For example, participation in UDDI can help an established B2B eCommerce player expand into new markets and services or allow any size company just entering the online space to accelerate toward a world-class business presence.

The UDDI specifications take advantage of World Wide Web Consortium (W3C) and Internet Engineering Task Force (IETF) standards such as Extensible Markup Language (XML), HTTP, and Domain Name System (DNS) protocols. Additionally, cross platform programming features are addressed by adopting early versions of the proposed Simple Object Access Protocol (SOAP) messaging specifications found at the W3C Web site.

The UDDI project participants are split up into two categories, working group members and advisory group members. The fourteen working group members gather input from the 300-plus advisory group members to develop the UDDI specification. After the third revision of the specification, the UDDI project will hand it over to a standards body for consideration as a standard.

More Than Just a Specification

The UDDI specification differs from other registry efforts by virtue of the substantial commitment from industry partners to use this technology and implement it in their core businesses today – ensuring the specification will truly solve the problems facing small and large businesses, marketplaces, and technology providers.

As a demonstration of this commitment, several companies (ie: SAP, IBM, HP, and Microsoft) are launching a jointly operated UDDI Business Registry on the web. The UDDI Business Registry is an implementation of the UDDI specification and will enable all businesses to leverage this effort in their Web services activities.

The UDDI initiative leverages industry standards such as HTTP, XML, SOAP, and other specifications, thus demonstrating the openness of the approach and platform-independent *commitment*. This foundation gives all businesses that register in the UDDI Business Registry a kind of ID card, a globally unique identifier for them as a business.

UDDI-based registries reach far beyond today's Internet business listings and search directories that provide specific, but limited value to an organization. Major benefits have been historically derived by the use of widely adopted standards in all industries and/or initiatives. UDDI-enabled businesses will realize unprecedented value from the rapid acceleration of service-based computing as a result of this global initiative.

Summary

In summary, UDDI-based registries will become the cornerstones of web services, making it easier to publish your preferred means of doing business, find relevant services and have them find you, and integrate with these services over the Internet. By automating these processes:

- Businesses will have a means to describe their services and business processes in a global, open environment on the Internet thus extending their reach.
- Potential trading partners will quickly and dynamically discover and interact with each other on the Internet via their preferred applications thus reducing time to market.
- The barriers to rapid participation in the global Internet economy will be removed for any business anywhere thus allowing them to fully participate in the new digital economy.
- Businesses will be able to organize their portfolio of services in a controlled environment.