THE VIRTUES ARCHITECTURE: A SOFTWARE INFRASTRUCTURE FOR BUSINESS-TO-BUSINESS E-COMMERCE

P. Nixon, V. Wade, S. Terzis, M. O'Connell, and S. Dobson
Computer Science Department, University of Dublin, Trinity College, Dublin 2, Ireland.
{ Paddy.Nixon, Vincent.Wade, Sotirios.Terzis, Marcus.Oconnell, Simon.Dobson}@cs.tcd.ie

Abstract: Most discussion on the current hot topic of e-commerce focuses either on the relationship between the customer and the supplier or on the security of these interactions. In this paper we present an architecture for building business-to-business e-commerce applications. The primary objective is to develop a lightweight infrastructure for building and maintaining collaborations from partners distributed across the Internet. Within this infrastructure, techniques have been developed to aid in the location and composition of services, manage the distributed workflow process and aid in maintaining contractual   obligations.

1.   INTRODUCTION

Modern business practices are undergoing a dramatic shift. Business imperatives such as accelerating product cycles and improved product targeting imply that both fundamental (long-term) and market-driven (short-term) collaborations are critical to a business' continued competitiveness. Moreover, downsizing and narrow niche markets suggest that collaboration has great potential for reducing costs and fostering interoperability of products and services within the global marketplace.

Although the Internet, and e-commerce in particular, promise to enhance the way we do business their most significant contribution to date has been in the one-to-one direct selling domain. However, e-commerce is much more than just direct sales. E-commerce is about harnessing technology to support every aspect of business. An area of particular interest is how supply chains, or business-to-business relationships, can be improved through the use of the Internet. Particularly, Internet technology promises to make the process of building such supply chains faster, more open and more profitable. A number of problems manifest themselves immediately when considering how to implement such supply chains over the Internet to realise these promises. The location of potential collaboration partners is complicated, as is the logistics of agreeing a suitable collaborative contract. Once formed large project consortia, and those with highly dynamic populations, are notoriously hard to administer and control. Finally there are important legal and commercial issues in the accessibility of information

between partners, both during the project and (equally importantly) when it has been formally completed

This paper starts by providing a description of business-to-business collaborations through a brief definition of dynamic virtual organisations. It then outlines the VIRTUES architecture and the assumptions that underpin it. Two aspects of the architecture will then be presented, namely: service location, and workflow management. The paper will then concludes with a discussion of the architecture and how it might contribute to solution of fundamental problems in business collaborations such as trust management and contract agreement and obligation.

2.   VIRTUAL ENTERPRISE MODEL

A virtual enterprise (or virtual organisation) is the essence of a business-to-business relationship and is defined as an association constructed from both administratively and geographically distributed business units or organisations. It is a set of legally independent performers of varying types who voluntarily co-operate to seize market opportunity. They are represented by at least one partner to the external world and they agree to produce a common output, e.g. a product or a service, based on a common understanding of their business rules and business processes. In general, in a virtual enterprise environment a set of business processes are shared according to well-defined contracts and agreements. Of key importance to the successful implementation of a virtual enterprise is an architecture that enables the integration, sharing and management of business processes located in different business domain boundaries.

For an organisation to be referred to as virtual, it needs to base its co-operation or rather sheer existence on the use and application of information technology (IT). Hence, the use of IT is a constitutive feature of the virtual organisation. This allows it to be differentiated from other types of networked organisations - The virtual organisation is a network organisation but, in addition to implementing various forms of co-operation, it makes a heavy and critical use of information technology. Hence, IT emerges as the primary integrator of the virtual corporation. Information technology transcending organisational boundaries spans companies together into an agile and re-configurable network of high efficiency and adaptability. Only recent developments in network computing and the Internet have made a truly global and efficient virtual organisation a viable idea.

However, the virtual corporation is not just a collection of partners, but a collaborative structure, and this amplifies its apparent lack of boundedness. This is because co-operation ties the collaborators together to such an extent that they are practically merged into one, though re-configurable, structure. Since each constituent realises only a special fraction of the value chain, on their own, constitutive parts are nothing. The whole situation is further amplified by the fact that virtual partners share their resources, infrastructure, personnel, research, information and knowledge.

Virtual organisations have two key structural characteristics: interdependence between the constituent operations, and distribution of responsibility between constituent operations. They are globally distributed, and exploit information and communication technologies to support their operation. Information systems allow

virtual organisations to monitor feedback and refine their configurations, allowing them to constantly evolve. Appel states that there are five key types of virtual organisation:

1. alliances of organisations
2. alliances of individuals
3. established decentralised companies
4. central companies seeking to adapt
5. single organisations

Moreover, virtual organisations have at least one of the following four characteristics: geographic separation, functional specialisation with separate reporting hierarchies, transitory membership driven by evolving needs over time, and separation of production across different time dimensions.

From the discussion presented above we can see that *dynamic collaboration* using IT infrastructures is central to enabling business-to-business e-commerce. It is our contention that this model will be central to the business organisation of the future and that current software infrastructures are inappropriate for, or are unable to support, such relationships.

3.   VIRTUES ARCHITECTURE

The VIRTUES system addresses the observation that no two businesses are going to be prepared to modify either the businesses processes or underlying technology for every electronic collaboration they undertake.
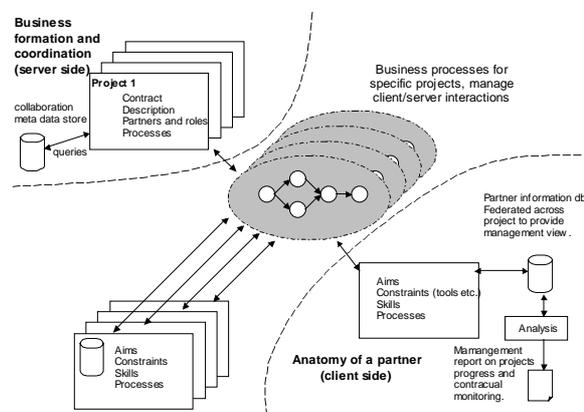

Figure 1: The VIRTUES architecture: management view

The VIRTUES architecture can be viewed from two perspectives: the management or global perspective and the systems level perspective. Figure 1 shows the management perspective that depicts a high level view of a set of collaborating companies (bottom right), managed by a coordinating partner (top left), with a business process (in the centre) per grouping facilitating the interactions.
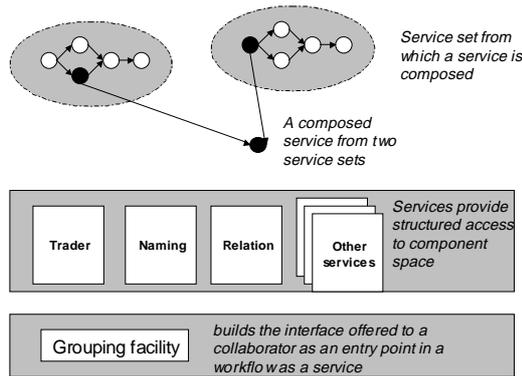
Figure 2: The VIRTUES architecture: Systems view

The basic assumption of this management view is that each collaboration, or supply chain, is sufficiently dynamic that collaborators can be added and removed by the coordinator and that interactions between collaborators are governed by the appropriate business process. The systems level view is depicted in figure 2. Here is where the problems of technology integration are handled. At this level all collaborators are viewed as a set of services embodied in software components. The location and integration of the components is performed via an enhanced trading service. The trading service is support by facilities for locating a named component (naming service) and a facility that maintains a graph of relationships (such as `requires`) about named components. The naming and relationship services follow OMG definitions. Moreover, the realisation of the business process is also made this level as a distributed workflow support by a workflow engine. Below we introduce two aspects of the system which are unique to VIRTUES; namely the component and service location mechanism and the workflow system.

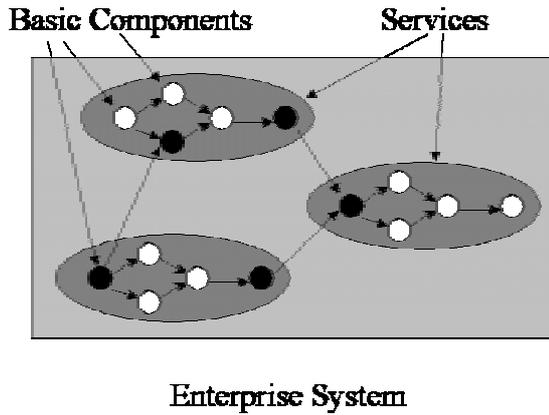## 3.1 Service location and composition

Figure 3: Level of abstraction for enterprise system decomposition.

To take the component-oriented view of enterprise system development, presented above, requires the decomposition of current enterprise systems into a number of services and components. This system decomposition is done according to different levels of abstraction (see Figure 3). At the lowest level the system is decomposed to a number of basic components. Some of these components grouped together form functionally cohesive entities called services. At the highest level, these services are combined together to form the whole enterprise system. In the development of a VES the participating systems could contribute either basic level components or services. So, the facility should support component location (aka company negotiation) and composition (aka company formation) at both levels. The approach we propose does this via a trading facility.

The component trading facility supports the development of VES in the following way (see Figure 4). At the start there are a number of existing enterprise systems. Each of these systems consists of a number of connected components (both basic level components and services from) and a semantic trader (the black circle inside each system). The creation of the VES (the system in the middle) will be driven by a new component trader, which is formed by composing the existing ones. Then, the new trader is used to select the components that will form the VES while the composition will be supported by a wrapping service (explained below) associated with the trader creates the new configuration. Some times the existing systems might not provide all the necessary components in which case the component trader will try to retrieve missing components from a worldwide pool of available ones.

- Enterprise Systems
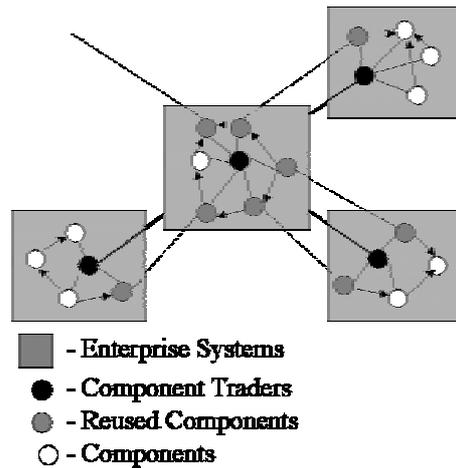- Component Traders
- Reused Components
- Components

Figure 4: Use of Component Trading Facility for VES development.

This sort of trading cannot be performed by current industry standard trading service such as the CORBA trader. Therefore, we designed a semantically enhanced trader that performs a fundamentally different type of trading: its foundation is answering questions of the type, "*Find me a component/service **that does the following**". So, the focus shifts from the appearance (syntax) to the functionality (semantics) of the components. The process of trading is governed wholly from the management view via the coordinating partner(s) as part of the start-up phase of a supply chain.

3.2   Distributed Workflow management

As already stated the realisation of the business process is achieved through the workflow engine and accompanying technologies. The core difference that is dicatated by the context is that such a workflow engine must be distributed. The engine has been developed to support the integration of distributed information systems.

The engine itself consists of a scheduler, which accepts management requests and initiates instances of these management processes. The scheduler uses a Knowledge server to interrogate the management process rule base and determines the next activity to be enacted. The scheduler is implemented as a multi threaded process in order to deal with concurrent management requests. When the scheduler initiates work, this work is logged within a Workflow Information Server (WIS). This WIS server maintains the state of all management process instances (i.e. all instances of management requests currently being executed within the management system). Once the next activity to be enacted as part of a management process has been identified, the scheduler passes this information to the workflow dispatcher.

The dispatcher is responsible for the invocation of the appropriate management component, which supports this activity. Figure 5 depicts the engine and illustrates the components in the system.
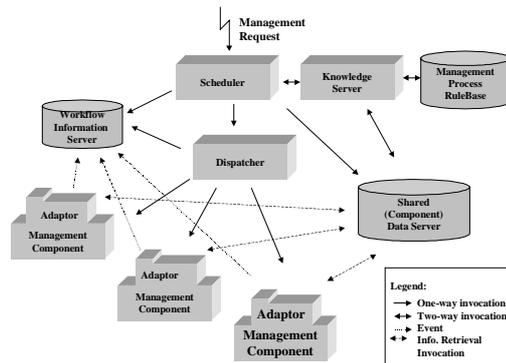
Figure 5: The VIRTUES workflow management architecture.

The dispatcher is likewise multithreaded to support concurrent component invocations. Typically before a component can be invoked, some input parameters have to be retrieved. Such parameters may be configuration information, or may be outputs from the execution of other component invocations. The dispatcher could potentially become congested if it must perform this information gathering as well as carry out concurrent component invocations. Also, the implementation of the dispatcher may become very complex if it has to know or interpret the information (parametric) requirements of each component.

For this reason, component adaptors were developed which interface the workflow engine to the components. The adaptor source code is over 50% generic as the interface to the workflow engine is standardised across all workflow adaptors and only workflow control data is passed between the engine and adaptor(s). The management component specific part of the adaptor is responsible for retrieving the information required to invoke a management component. This information is stored in the Shared (Component) Data Server, the interface to which is again common for all adaptors. The adaptor is also responsible for placing any resultant information, which is required to be shared, into the Shared Data Server. A 'wrapper' object, either remote or running in the virtual memory space of the adaptor, performs the actual interaction with the management component.

The adaptor lets the workflow engine know that specific management activities have been completed, by sending events or one way asynchronous calls to the Workflow Information Server (WIS). The WIS has a number of registered receivers, which require to be notified of such completions. These include but are not necessarily limited to the scheduler and dispatcher.

The use of asynchronous invocations (or events) between the adaptor, WIS, Scheduler, and Dispatcher allows greater degree of concurrency, less chance of activity blocking and more flexible integration of the workflow engine itself.

An important aspect of the engine is that it can be federated and so co-operate to support business processes. This could be performed where there are local grouping of components (provide by the trader), but where these groups themselves

are geographical dispersed by a geographic distances, have congested or poor network connectivity or are under the control of separate administrations.

3.3 Other services

Additional services that are supported in the architecture but not reported here are: service mobility/migration, contract description and monitoring, along with a programming language system that facilitates all of the above.

4. CONCLUSIONS

We have presented a novel architecture for providing a lightweight software infrastructure for business-to-business e-commerce. Two of the novel features of this architecture have been presented in some detail: the enhanced trading service and the workflow engine. Both of these services are departures from the usual approaches, as they have to support radically different requirements, namely: highly dynamic collaborations requiring services to be located, composed, and put into operation. The new trading service achieves this through the addition of semantic trading, providing a mechanism for component and service location and composition. The workflow system implements the necessary decentralised engine necessary for implementing and integrating heterogeneous business processes.

In conclusion, the integration of trading, grouping, composition and workflow at an enterprise level, along with additional services, provides an appealing approach to supply-chain construction.

However it raises some deeper questions of trust management and contract obligation. Now that companies can *collaborate* on the internet how can we ensure they are who they say they are? How can we force them to meet their obligations through contracting? How do we audit the overall process and how do we manage the ownership of product and data after the lifetime of the collaboration.

A hint to answer for the more tangible problems lies in the nature of our encapsulation of service as software interface that can be (although not easily) rigorously specified and verified and hence can be expressed as a machine readable and auditable contract. The intangibles, such as building trust relationships, are harder to fathom and will probably be solved by brokers who vouch for the track record of a company. We are actively pursuing the answers to these, and other, questions as follow research within the VIRTUES architecture.

ACKNOWLEDGEMENTS

REFERENCES

Terzis. S, Dobson S.A., Wade V., **Nixon P.A.**, & Fuller J. *Building the next generation groupware*, in proceedings of International Conference on Enterprise Information Systems, Editors: Joaquim Filipe and Jose Cordeiro, pp 525-532. 1999. *Best paper award.*

Baker, S., Cahill, V. and Nixon, P.A. 1997, *Bridging Boundaries: Corba in perspective*, IEEE Internet Computing, vol. 1, no 5, pp. 52 – 57

Vaggelis Ouzounis, Volker Tschammer. *Integration of Electronic Commerce Business Processes* in *Virtual Enterprises*. EMMSEC98, European Multimedia, Microprocessor Systems and Electronic Commerce. 28-30 September 1998.

Tom Digre. *Business Object Component Architecture.* IEEE Software September/October 1998.

Wolfgang Appel. *Towards the theory of Virtual Organisations: A description of their formation and figure.* virtual-organization.net Newsletter Vol. 2, No. 2.

Marie-Claude Boudreau, Karen D. Loch, Daniel Robey, Detmar Straud. *Going Global: Using Information Technology to Advance the Competitiveness of the Virtual Transnational Organization.* Academy of Management Executive, Vol. 12, No. 4. 1998.

S.Terzis and Nixon P.A, *Semantic Trading: Tackling Interoperability Problems during System Integration*, Editors: A. Vallecillo, J. Hernandez, J.M. Troya, Object Interoperability, Selected Papers from ECOOP 99 workshop on Object Interoperability, Universidad de Malaga, Depto. Lenguajes y C. de la Computacion, ISBN 84-699-0520-1