Workflow Interoperability - Enabling E-Commerce

Mike Anderson MSc, BSc, MBCS, CEng
International Computers Ltd

Rob Allen BSc
SNS/ASSURE Corp.

Abstract: *This paper describes the rationale for workflow interoperability in the context of electronic commerce and as a means of implementing value chains that operate across and between organisations. The paper outlines initiatives undertaken by the Workflow Management Coalition (WfMC) to promote this capability through the use of Standards. A status update on which vendors can offer this capability is provided together with an explanation of how the capability of particular vendors can be assessed.*

The Commercial Opportunities for workflow Interoperability

E-Commerce is a pervasive activity. With the growth of the internet, it has caught the popular imagination and has even penetrated political thought in ways that we have not seen for a generation. Everyone has a view or an understanding of what they perceive E-Commerce to be. However, because it is a phenomenon that builds on the contributions of many different underlying technologies, no two experts will be able to give the same precise definition of what they mean when they use the term. For the purposes of this paper, we might use a definition put forward by the Chase Manhattan Bank:

> *"The distribution of products, services and/or information that involves a **requestor**, a **delivery channel** and a **provider** using an accepted methodology and an acceptable assumption of business and technical risk"*
>
> Chase Manhattan Bank

This definition indicates that most commercial computing falls within the definition of E-Commerce. E-Commerce includes both the digital transformation of recognisable commercial activities and the creation of new business rules and roles for participants in emerging arenas. The delivery of business through E-Commerce involves the deployment of business processes for which workflow is an obvious supporting technology. If we look again at the definition above, we can see that it is likely that the delivery of goods and services through E-Commerce will necessarily involve the operation of business processes that run across and between organisations.

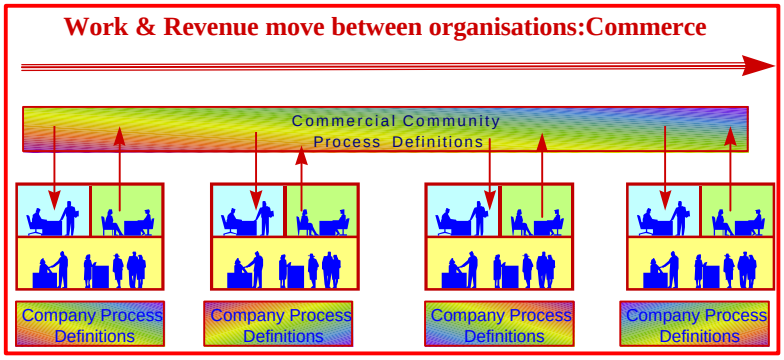

The implementation of value chains that run across and between organisations is not, in itself, a new idea. Electronic Document Interchange (EDI) has provided a reliable messaging regime to support inter-trading between consenting organisations for some time. Inter-trading is effected through the exchange of messages containing standard business objects (documents such as invoices, purchase orders or electronic funds) which are treated as input to the receiving organisation's IT systems. EDI regimes are well suited to support of secure, high volume, transactional inter-trading applications. Experience has however, shown them to be expensive to set up and somewhat inflexible once in operation.

What is Workflow Interoperability?

Business processes that operate within, across or between organisations in order to implement value chains that can be used to deliver E-Commerce transactions may be implemented using a set of workflow definitions that have been created to support discrete segments of the overall process. This scenario poses the question of how to avoid creating islands of automation in the operation of an end to end business process. The answer to this problem is workflow interoperability – the enabling of different workflow products to *"talk to each other"* by exchanging messages that effect process interoperation and integration to drive and manage the operation of the value chain. Workflow interoperability enables the owner of the value chain to have greater visibility and control over its performance and participants within the value chain benefit from flexibility, and improved control and visibility over the performance of the processes they operate and the processes with which they interoperate.

Supply Chain Management

Supply chain management is that subset of value chain engineering that deals with the supply of manufactured goods. It potentially addresses all significant activity that occurs from accepting an order through to the delivery of the finished product. Supply chains may require the active participation of wholesalers, hauliers, manufacturers, brokers and bankers as well as customers.

A supplier operating within a supply chain will typically operate four basic business processes:

- **Planning** which begins when the supplier's customer attempts to place an order. The planning process enables the supplier to reconcile when the customer wants the goods with when they can actually have them. It also deals with mundane matters such as production scheduling, inventory management and transport scheduling.

- **Sourcing** which deals with ordering of materials, delivery scheduling and processing of invoices.

- **Making** is the manufacturing process that results in finished goods

- **Delivering** which deals with how the finished goods move from the supplier to the next point of activity within the value chain.

In the past, manufacturers were the drivers of the value chain, controlling the pace of production and distribution of goods. Today, customers are much more demanding and in many industries have taken ownership of the value chain. In such a competitive climate, participants in the supply chain must constantly seek to improve their business processes to meet rising customer demands for more choice, more efficient production, speedier, more predictable delivery and even the ability to track progress on order fulfilment.

The following is a supply chain scenario that, in various forms, has been used as the basis of a number of interoperability experiments and demonstrations by WfMC member companies. The scenario describes the life history of an order fulfilment for a product (lets say a new sofa) from a retailer (lets say a mail order company) which places orders with manufacturers rather than holding stock. The retailer must therefore make arrangements with the manufacturer to produce the goods (a new sofa) and with a transport company to deliver the goods to the customer.

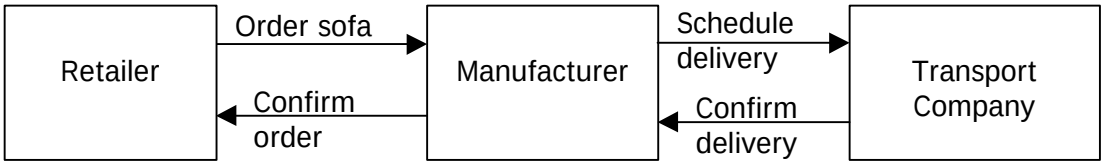

Figure 1: The order fulfilment scenario

Each participant in the supply chain uses a *workflow engine* to manage the processes it operates to deliver its contribution towards order fulfilment. A *workflow engine* is a software product used to manage the routing and scheduling of task oriented work according to a pre-ordained *process definition*. When a new process starts (say a customer places an order), the workflow engine reads the appropriate process definition and starts the first defined *activity* which, in the case of the retailer, may be to display an order form for completion.

Each new process that is started on a workflow engine is known as a *process instance*. A *process instance* is a defined thread of activity that is being *enacted* (managed) by a workflow engine. Most workflow engines can report on the current status of a given process instance.

When an activity is completed (say the order form has been filled in), the workflow engine uses the process definition to decide what to do next. The next activity might be for an order to be placed with the manufacturer to make a new sofa. To place the order, our scenario says that the retailer's workflow engine makes a request to the manufacturer's workflow engine to start enactment of a process instance that will manage activities associated with the:

1. production scheduling
2. delivery scheduling
3. manufacture
4. proof of despatch
5. billing

relating to the production of the sofa. To make this happen, the retailer's workflow engine will send a message to the manufacturer's workflow engine requesting that it start an instance of its order fulfilment process. Elements of process relevant data that must be included in the message in order to *instantiate* (pass values for process variables) the process model will include:

- Order number
- The product code (identifying the product as a sofa)
- The colour of the sofa to be made
- The customer's name
- The delivery address
- The date offered to the customer for delivery

The first step in the manufacturer's defined process is to invoke an MRP application to schedule production of the sofa. Once the date the sofa will be ready for despatch has been established (returned to the workflow engine by the MRP application), the process moves

onto the next step which is to book delivery with the transport company. This is done in the same way that the order was placed with the manufacturer. A message is sent to the transport company's workflow engine requesting that it start an instance of its book delivery process and providing the following elements of workflow relevant data to instantiate the process definition:

- Where to pick up from
- What to pick up
- Value of goods to be shipped
- Delivery address
- Name of recipient
- Date for collection
- Collection time
- Date for delivery
- Delivery time

The transport company will schedule the delivery and return the confirmed delivery date to the manufacturer's workflow engine.

The manufacturer's workflow engine will return confirmation of the delivery date to the retailer's workflow engine, which will then send an invoice to the end customer stating details of what has been ordered and when it will be delivered.

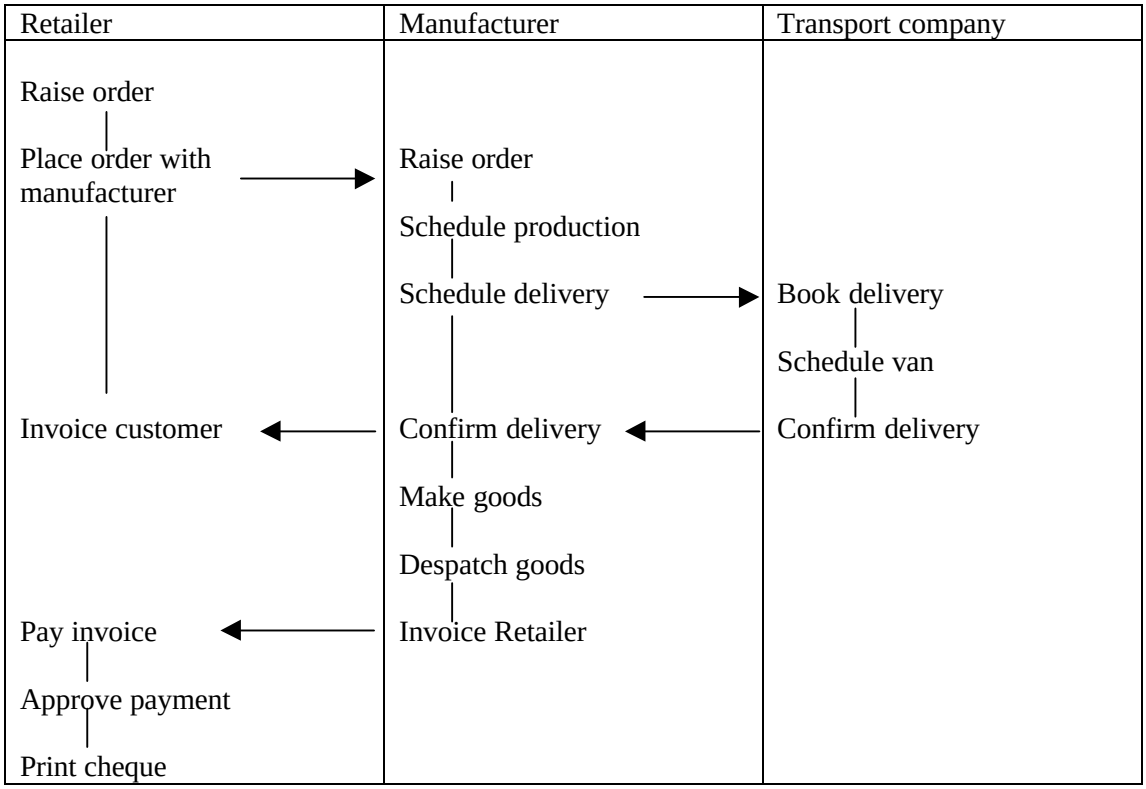

Figure 2: The order fulfilment process flow diagram

The manufacturer and transport company processes are modelled as *nested sub-processes*, i.e. they provide information back to their parent process instance enacted on the workflow engine that requested their invocation (see the process flow diagram for order fulfilment above).

An extension to the above scenario illustrates a second class of workflow interoperability based on the use of *chained sub-processes*. Having produced the sofa and despatched it to the customer the manufacturer raises an invoice with the retailer requesting payment. This is done electronically by the manufacturer's workflow engine requesting the retailer's workflow engine to start an instance of its pay invoice process and providing it with the following information:

- Invoicing company
- Amount
- Item purchased
- Order number
- Despatch date

In order to instantiate the process definition. The first step in the pay invoice process is an approval step. Assuming the payment is approved, a cheque will be printed and sent to the manufacturer by post. The manufacturer's process ends once payment has been requested. A *chained sub-process* is one in which the parent process starts a process instance on another workflow engine, but takes no further interest in its progress.

It is our expectation that *chained sub-processes* will be used more within organisations where they will be treated as members of a set of trusted processes deployed to support operation of the business. *Nested sub-processes* are more likely to be used in those business contexts where rendezvous and co-ordination are material to the operation of the business process or where it is important that the parent process track or maintain operational control over the operation of the sub-process.

For all of this to be possible, there will need to be a commercial agreement in place between each pair of companies that inter-trade in the value chain. The companies will also have established a technical agreement (an *interoperability contract*) describing the way in which interoperability is to be effected. The technical agreement will cover aspects such as:

1. which workflow engines within one company are visible to/capable of interoperating with which workflow engines in the other company
2. which workflow definitions can be enacted within one company at the behest of workflow engines in the other company
3. the transport binding supported (e.g. MIME[1], jFlow, SWAP,…)
4. for each workflow definition identified in the contract:
    - values that must be supplied or can be returned
    - for each traded (shared) element of *workflow relevant data*
        - access rights (whether the element is readable/writable)
        - value constraints (minimum/maximum values, number of permitted updates/accesses)
    - outcomes/outputs/returned elements of workflow relevant data
    - audit data policy
    - change control policy
5. security policy and implementation
    - authentication

[1] The Abstract Specification was produced in recognition that fashion and technology change. It is probable that different companies or industry sectors will elect to use particular transport technologies such as MIME or CORBA for effecting workflow interoperability. The principles that underpin what they are doing will however, remain the same. The WfMC has currently defined MIME and CORBA bindings and will bring forward other bindings (such as SWAP) as the needs of industry change or new needs emerge.

- support for/ policy on non-repudiation
- shared key cryptography & key management
- handling security breaches

6. exception handling/recovery protocols & transactional behaviour

Individual interoperability contracts will have a unique identifier, determined by the organisations trading across the service boundary, which is used to support authentication mechanisms.

Where we are today (and how we got here)

When the Workflow Management Coalition (WfMC) began its work to define standards for the workflow industry sector in 1994, achieving workflow interoperability was been seen as a key objective. Industry analysts and commentators identify workflow interoperability as a key deliverable from the Coalition. Four years after the U.S. Department of Defence sponsored an initial study to kick start the WfMC work in this area, a lot of progress has been made.

- June 1996 – First WfMC Workflow Interoperability demonstration at Workflow Canada 1996 in Toronto
- November 1996 – WfMC Workflow Interoperability demonstration at GIGA Show in Amsterdam
- November 1996 – WfMC Workflow Interoperability – Abstract Specification outlining the theory behind effecting workflow interoperability published.
- November 1996 – WfMC Workflow Interoperability – MIME Binding published. The MIME Binding defines a messaging protocol to be implemented using MIME constructs to effect the style of interoperability described in the Abstract Specification.
- May 1997 – WfMC Workflow Interoperability Demonstration at LOMA ’97 in Florida.
- November 1997 – WfMC begins work to define Interoperability Proving Framework as a mechanism for proving conformance.
- May 1998 – Joint WfMC/DMA Interoperability Demonstration at AIIM ’98 in Anaheim.
- October 1998 - WfMC publish revised draft of the MIME Binding based on feedback from demonstrations.
- October 1998 – WfMC issues the first Interoperability Challenge.

The demonstrations staged between June 1996 and May 1998 were used as a vehicle to validate and provide feedback on the Interoperability Specifications. The revised version of the MIME Binding (version 1.1) was published in October 1998 and is the basis on which the Coalition expects to see vendors bring product to market.

The Interoperability Challenge

In October 1998, following publication of version 1.1 of the MIME Binding, the WfMC issued its first interoperability challenge to the workflow vendor community in order to encourage vendors to bring product into the market place during 1999. To meet the challenge, a vendor must:

1. demonstrate that its product can complete the tests described in the Interoperability Proving Framework when interoperating with a product from another vendor
2. publish the conformance matrix which shows the test results
3. give a public statement of when it would expect to bring product to market.

The Interoperability Challenge will be run twice during 1999. The first testing session has been scheduled for March, allowing product announcements to be made at the AIIM Show in Atlanta in April 1999. The second testing session will be run during the Summer of 1999, allowing companies that have already committed their development resources to existing product development plans and release dates to also bring forward proven product this year.

Participants

First Wave (March 99)

	Contact Summary	Contact Details (phone Email)
DST		
FileNET		
Staffware		

Second Wave (Summer 99)

	Contact Summary	Contact Details (phone Email)
Computron		
CSE		
Documentum		
Fujitsu		
Hitachi		
K+V Van Alphen Automatisieringsdiensten		
SNS/Assure Corp	Rob Allen	(+44) 181 261 8205 rob_allen@sns.ca

Challenge deliverables

The work required to meet The Challenge involves:

1. Implementing 11 operations as specified in version 1.1 of the MIME Binding
2. Writing workflow definitions that will implement the Interoperability Proving Framework tests for each workflow engine
3. Proving that the tests work in house using two instances for each product
4. Running the tests against other vendor's products during Challenge Week over the internet
5. Documenting and publishing the test results.
6. A public Product Availability announcement. This should state when and in what form the product will be released. The Coalition encourages its members to release product in 1999, but the only prerequisite for participating in the Challenge is that a release date be given.

Assessing Interoperability

The WfMC Interoperability Standards are designed to allow users of workflow products to implement processes that flow across organisational and technological barriers. A difficulty faced by the authors of the Standard is that different workflow engines are founded on different conceptual models and have different behavioural characteristics and capabilities that affect the way in which they can support interoperability with other workflow engines.

Recognising the need to be able to assess the meaning of conformance statements that might be made by vendors, the W*f*MC Working Group 4 (Workflow Engine Interoperability) produced an Interoperability Proving Framework document. This work was done against a background of so-called conformance claims made by a number of companies that are not members of the Coalition. The document sets out a comprehensive series of tests designed to show the behaviour and ability to interoperate of two workflow engines that have implemented the MIME Binding specification.

The framework does not address how a given workflow engine is told that it needs to interoperate with some other designated engine. Rather, this is assumed and the framework addresses what must follow as a consequence. The reason for this is that in supply chains, there could be any number of participating engines. It is a simple fact that the design and performance of any two workflow engines will not be the same, consequently it is important to define interoperability externally to the engine.

Each interoperability capability is described in the form of a question to which there is expected to be a yes/no answer. The framework defines an *Interoperability Compatibility Matrix* (ICM) which can be used to assess the ability of specific workflow engines to interoperate using the defined MIME binding. A number of test scripts are provided with the resulting entries in the ICM given as examples.

Variations in workflow vendors' implementations of the Standard arise in areas such as:

- support for authorisation and security policies
- support for the collection of audit data
- support for different classes of process termination
- support for event notification

To allow some vendors to implement the Standard without having to completely re-engineer their products, the Coalition has agreed the following conformance levels that can be achieved:

1. *Chained sub-process interoperability*
2. *Nested sub-process interoperability*
3. *Process Management through interoperability*

The message protocols have been subsetted in line with these conformance levels and a framework of tests devised to promote accurate assessment of two workflow engines to effect interoperability at each level. The test results are tabulated in an *interoperability compatibility matrix* (see below), from which a potential user can understand the constraints or degrees of freedom that they incur by using particular product sets.

All of the tests are symmetrical (the functionality they are designed to evaluate must be tested in both directions). Thus, in the Interoperability Compatibility Matrix (see below) each test is enumerated as follows:

interoperability level . test . initiating workflow engine

In the example given below, the matrix shows that using two designated workflow products, the *WonderFlow* engine will maintain an audit trail of its view of the interoperations that occur, but this cannot be reflected for those parts of the process managed by the *SuperFlow US* engine. It also shows that there is no support for authorisation checking from the *SuperFlow US* engine (test 5 and test 10 require that the responding workflow engine is able to validate the authorisation of the requesting workflow engine). This information may be used as grounds for determining product selection or to inform process design. A user or systems integrator assessing the products using the Interoperability Proving Framework is empowered to make their own decisions.

Test	*WonderFlow*	Audit data recorded	*SuperFlow US*	Audit data recorded
1.1	Y	Y	Y	N
1.2/1	Y	Y	Y	N
1.2/2	Y	Y	Y	N
1.3/1	Y	Y	Y	N
1.3/2	Y	Y	Y	N
1.4/1	Y	Y	Y	N
1.4/2	Y	Y	Y	N
1.5/1	Not tested		Not supported	
1.5/2	Y	Y	Y	N
1.6/1	Y	Y	Y	N
1.6/2	Y	Y	Y	N
1.7/1	Y	Y	Y	N
1.7/2	Y	Y	Y	N
1.8/1	Y	Y	Y	N
1.8/2	Y	Y	Y	N
1.9/1	Y	Y	Y	N
1.9/2	Y	Y	N	N
1.10/1	Not tested		Not supported	
1.10/2	Y	Y	Y	N
1.11/1	Y	Y	Y	N
1.11/2	Y	Y	Y	N
1.12/1	Y	Y	Y	N
1.12/2	Y	Y	Y	N
1.13/1	Y	Y	Y	N
1.13/2	Y	Y	Y	N
1.14/1	Y	Y	Y	N
1.14/2	Y	Y	Y	N

Sample Interoperability Compatibility Matrix for 2 workflow engines claiming support for implementing *Chained sub-process interoperability (Conformance level 1)*

Interoperability compatibility matrices for workflow products that have been tested against each other under the terms of the Interoperability Challenge will be posted on the WfMC public web page (http://www.wfmc.org).

Summary
This paper proposes the ideas that E-Commerce is an activity that is important to the global economy, and that the W*f*MC Interoperability Specification facilitates the growth of E-Commerce.

The WfMC Interoperability Specification is an Industry Standard agreed by all the major workflow vendors. The MIME Binding to the Interoperability Standard provides the basis for all conforming products to communicate via e-mail that can be received by another workflow engine and used to initiate a new business process. The interoperability is bi-directional. All conforming products are expected to be able to send, receive and respond to these e-mails.

The Interoperability Standard enables provision of automated business processes that span many organisations to operate integrated supply chains. It is cheaper to operate commercial transactions automatically, and this is why EDI has been so popular. Workflow

interoperability allows organisations greater flexibility in how they meet their commitments as participants in a supply chain.

A number of workflow vendors have demonstrated compliance during March 1999 and announced their success at the AIIM show in Atlanta in April. Others will prove compliance during the Summer of 1999. A number of workflow vendors will release products that will provide the capability to inter-operate during 1999.

Organisations within industry sectors will discuss how they can best regulate themselves so that they can gain the benefits of automation. The participants in a value chain will define the business processes that collectively span the contributions made by all organisations. Such contributions will, by their very nature, be modular, replicable and replaceable. Value chains will develop that are long but fairly simple. Value chains with many sub-processes are also anticipated, and these will be deep rather than long.

Value chain members will work to develop specific process definitions that relate to the contribution(s) of each participant. This activity will be underpinned by trading agreements that are described as both commercial and interoperability contracts. Testing of the processes composing a value chain may be incremental (as automation migrates across the value chain). Testing must ensure that the integrity of the end-to-end process is preserved, that the process gives the desired result, and when an instance falls outside agreed tolerances, that the exception handling meets the commercial needs of all relevant interests.

Eventually, this and similar technologies providing commercial inter-operability will be very wide spread. Current indications are that global manufacturers will be the first to implement sizeable commercial communities. For example, the major motor manufacturers have substantial numbers of suppliers. These in turn use many component manufacturers that depend upon sub-contractors. The benefits of supply chain management are very clear in this environment. This means that the major companies will be encouraging smaller companies to co-operate throughout supply chains.

References

Anderson M.J., *Workflow Interoperability – what's in it for the user?,* Document World, May 1997

Hollingsworth D., *Workflow?A Model for Integration,* ICL Technical Journal,*Volume* 12, Issue 2

WfMC-TC-1012 *Workflow Standard – Interoperability Abstract Specification*, 1.0, Workflow Management Coalition, November 1996

WfMC-TC-1018 *Workflow Standard – Interoperability MIME Binding*, 1.1, Workflow Management Coalition, October 1998

WfMC-TC-1021 *Interoperability Proving Framework*, 1.0, Workflow Management Coalition, April 1999.

Further Information

Association for Information and Image Management International http://www.aiim.org

Black Forest Group http://www.penfield-gill.com

Data Interchange Standards Association	http://www.disa.org
Supply Chain Council Inc.	http://www.supplychainmgt.com
Workflow Management Coalition	http://www.wfmc.org