Workflow - A Model for Integration

David Hollingsworth

Principal Architect, Skill Centre, ICL Enterprises, Windsor, UK

Abstract

This article reviews the nature of a workflow system from a systems integration perspective, focusing on points of interaction between the workflow control software and other system components, such as process design tools, legacy applications and messaging infrastructure. The characteristics of the underlying business model and its representation to the workflow system are also discussed, including requirements for business processes to span organisational boundaries. The complexity of systems integration is identified as a major constraint on effective exploitation and indicative of the need for standards to support more effective product usage and interoperability. The article draws on the Author's experience in developing workflow related standards and concludes with an assessment of their potential impact, particularly on opportunities for their use in electronic commerce.

1. Introduction

Workflow is often seen as a key integration technology, bringing together business processes with the information to support them, and linking legacy and desktop applications into a flexible and adaptable distributed infrastructure. The external image of such systems can be deceptively simple, based on the notion that once the business process is defined, its automation merely requires the integration of a few simple tools.

According to the Workflow Management Coalition[1] [1], workflow represents "*the automation of a business process, in whole or part, during which documents, information or tasks are passed from one participant to another for action, according to a set of procedural rules"*.

Whilst not explicitly stated in the above definition, a key motivation for the deployment of workflow technology is that it should provide flexibility for the business process to evolve with minimum re-engineering. This is a concept simply captured within Open*framework* [2] *as "potential for change*".

Workflow technology typically achieves this by enforcing separation between:

- the definition of the various activities within the business process and their data requirements

[1] The Workflow Management Coalition (WfMC) is a non-profit making consortium of vendors, users, analysts and academia, with the goal of developing standards for workflow systems operation and promoting knowledge of the technology within the industry.

- the business rules governing the flow of control between activities within the process
- the roles and responsibilities associated with the work undertaken within the process activities
- an underlying organisational model, which relates roles and responsibilities to the actual work performers

In theory any aspect can change independently by simple amendment of the relevant control parameters, without affecting the ongoing operation of any other aspects of the process.

Despite this apparent utopia, the reality in many workflow systems implementations has been much more earthly - substantial systems integration issues to be faced in bringing together the component systems elements, lack of interoperability between different systems, major cultural and organisational issues to be resolved in the introduction of new working practices, and so on. Once operational, many systems prove less adaptive than expected to the future organisational or business changes .

2. Integration Requirements

Reliable statistics from within the industry are not always easy to find. However a recent market survey undertaken by the Workflow Management Coalition indicates that for virtually all workflow systems, integration with other industry software is vital - and a major cost component of implementation. At the time of writing the full survey results are still being collated but preliminary findings are as follows:

Integration Requirements	% of Respondents
World Wide Web	89%
Java Applications	75%
Legacy Applications	73%
CORBA based infrastructure	69%
Security Services	66%

Other technologies frequently mentioned included Business Process Modelling tools, Document Management and Imaging systems.

Informal estimates have indicated a ratio between workflow software implementation and overall project integration costs of between 1:5 (for ad-hoc office based systems) to 1:7 or more (for highly structured production workflow applications). Even in an industry where integration is increasingly the major cost component in the introduction of new technology, these are high figures.

This informal view is supported by a recent study from Ovum Group [3], which shows workflow vendor revenues split between product and service currently in the ratio of 1:4.5, with a fall projected by end 2000 to 1:3. This is consistent with an

increasing degree of standardisation, de jure or de facto, and of product consolidation leading to simpler integration.

The following sections consider two factors, the fragmented emergence of workflow within the market and the technical complexity of product interfaces, which have contributed to this cost.

3. The Evolution of Workflow Technology

One of the reasons for the complexity of the systems integration task is the fragmented way in which workflow technology has developed in the market.

3.1 Workflow - the first phase market

Software to control process operations is not a new concept. Many types of product in the IT market have supported certain aspects of workflow functionality for a number of years, although often embedded within other, related products rather than as a technology in its own right.

Image Processing

Workflow has been closely associated with image systems and many image systems have some workflow capability built-in. Once paper based information has been captured electronically as image data, it is often required to be passed between a number of different participants for different purposes within the (previously paper based) process.

Document Management

Increasingly, the management of electronic documents has included facilities for routing documents (in whole or part) between individuals and repositories, for example to facilitate shared authoring or filing services. Standardisation within the document management area has already recognised the requirement for extensions into workflow[2].

Electronic Mail & Directories

Electronic mail provides facilities for distributing information to individuals; associated directories provide a means of recording information about user attributes, such as organisation roles or other attributes relating to business procedures. Through the addition of routing mechanisms to define a sequence of recipients for a mail item electronic mail systems have themselves been progressing towards workflow functionality.

[2] The Open Document Management Association (ODMA) first identified an API for simple workflow functionality in 1995. More recently the DMA (Document Management Association, representing major vendors of document management software) has entered discussions with the WfMC to address the integration of workflow and document management standards.

Groupware Applications

Groupware applications are designed to support and improve interactions between groups of individuals. Initially many of these applications supported improvements in group working via informal processes, accessing group bulletin boards or diary/scheduling applications on an ad-hoc basis. As the scope of such applications has spread towards more formal business processes there has been an increasing move to incorporate workflow into work-group software[3].

Project Support Software

Software to handle complex IT project developments has often provided a form of workflow functionality within the project environment, for the sequencing and routing of development tasks between individuals and routing information between individuals to support these tasks.

Transactional Workflow

As traditional TP applications have become more distributed in nature some have moved to fully distribute transactional tasks to desktop environments. In parallel workflow vendors have been adding transactional characteristics to workflow systems, particularly in the area of task commitment and recovery co-ordination. In these situations there is an increasing degree of overlap between the two technologies.

BPR and Structured System Design Tools

Whilst workflow has been emerging as a fragmented technology, Business Process Re-engineering tools have appeared in significant numbers. These provide IT based support for analysing, defining and modelling the business processes of an organisation and the potential effects of change in such processes or organisations. The use of such products forms a natural pre-cursor to workflow implementation.

In summary, there are now many products in the market providing workflow capability[4]. Such products are often derivatives of products from other market areas, incorporating elements of workflow technology in an incompatible manner, making integration costly and negating the "potential for change" factor.

3.2 Workflow - the second phase market

The GIGA group recently presented an interesting analysis of the development of workflow technology and concluded that the industry is entering the second phase of automation [4]. Most of the first phase automation projects have been at departmental or workgroup level, with relatively little co-ordination[5]. The continuing business pressures of globalisation, contracting and electronic trading are leading organisations to reassess their business processes at enterprise level with ever increasing frequency.

---

[3] Examples of this include the integration of Lotus Notes product with several workflow packages and the introduction of Fujitsu/ICL's TeamWareFlow as the workflow component within Team Office.

[4] Of the WfMC membership, there are approximately 100 different organisations which categorise themselves as product vendors.

[5] At a recent conference a large multinational organisation indicated that of the 8 workflow applications implemented to date all were incompatible in terms of interoperability and use of common infrastructure components.

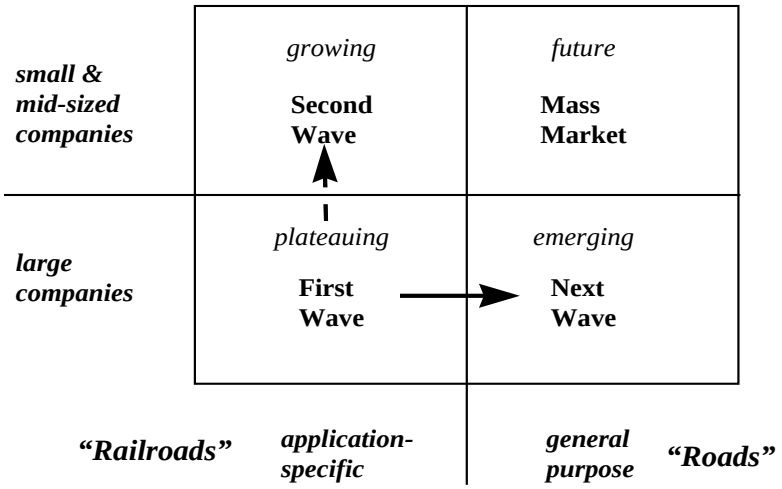

Figure 1 - Workflow & the shift to IT Infrastructure
(Giga Information Group, adapted from CAP Ventura)

The GIGA view is that "*second phase, messaging based enterprise-wide workflow will be dominant in 1-2 years time*". This will see workflow positioned as general purpose middleware across the enterprise. Electronic trading between organisations will increasingly push workflow into smaller and medium sized organisations, leading towards ubiquity of use.

This prognosis, however, depends upon the consolidation of the industry around a cohesive set of standards to support integration and interoperability, considered in Section 6.

4. Workflow and Software Integration

Integration complexity arises from the requirements of most workflow systems to interact with numerous other software components, ranging from standard desktop tools such as forms, spreadsheets and word-processors, to server applications such as document repositories and legacy applications, often based upon TP technology.

A key aspect of many workflow system is the incorporation of an organisational model, enabling workflow procedures to be defined relative to organisational roles and responsibilities. These may be separately maintained, for example by means of a directory subsystem, with associated role privileges.

Workflow systems may also require integration with process definition and modelling tools so that a proposed system can be fully specified and simulated prior to introduction.

Finally, as with any distributed application, integration with the underlying infrastructure (Electronic Mail, Object Request Broker domains, etc) is a further requirement.

The following schematic gives some indication of the potential components and points of integration of a typical workflow system.

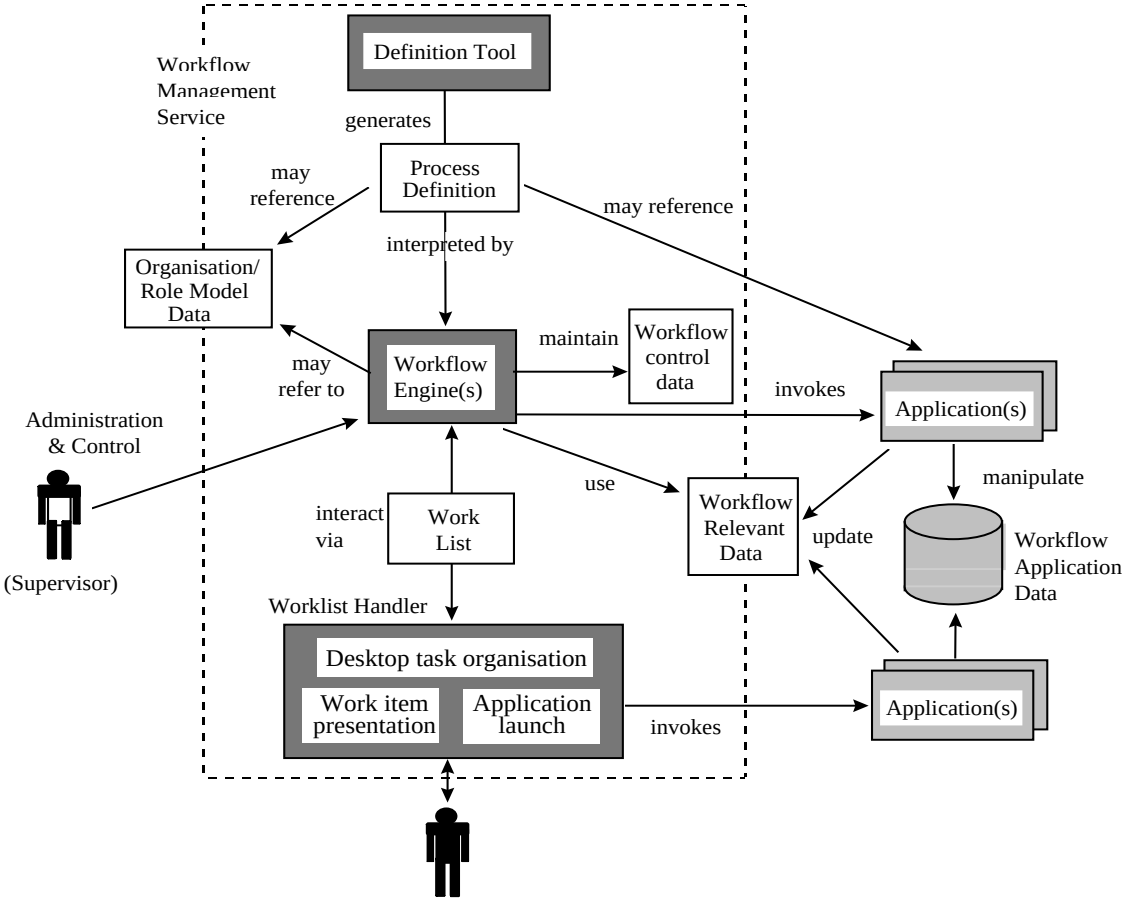

Figure 2 - Workflow System components

Several different systems construction paradigms exist within the industry. Common models include:

- Object based, using CORBA as the main distribution mechanism
- Electronic mail based with autonomous desktop environments
- Centralised workflow engine with tightly coupled desktop task management
- Document-centric with shared repository

The boundaries of the workflow management software are often unclear, for example some vendors include a directory component or interfaces to access legacy systems, others see this as part of the system integration task.

5. The Characteristics of the Business Process

Workflow is a process centred technology. To quote from Koulopoulos [5], of the Delphi Group (a Boston based workflow consulting group):
"Workflow emphasises the importance of the process, which acts as a container for the information. … This is a process-centred model, as opposed to an information-centred model."

5.1 The nature of the business model

Although the majority of workflow systems have tended to automate administrative processes (the so called “paper factory”) in an essentially human environment, there are often certain activities which are wholly automated by software components. In the manufacturing or process industries many activities are fully automated with little or no human involvement. In this context the specification of the work “performer” for a particular activity must incorporate the concept of machine automata.

Various characteristics of a business process need to be considered:

5.1.1 Responsibilities

“Ownership” of a business process is often an alien concept, but once an electronic representation is achieved, this becomes an important attribute of the process, if only to determine who has permission to modify the process and under what circumstances.

The realignment of business thinking from organisation to process marks a major shift in organisational culture, likened by Giga Group to the “*dismantling of the industrial age*”[4].

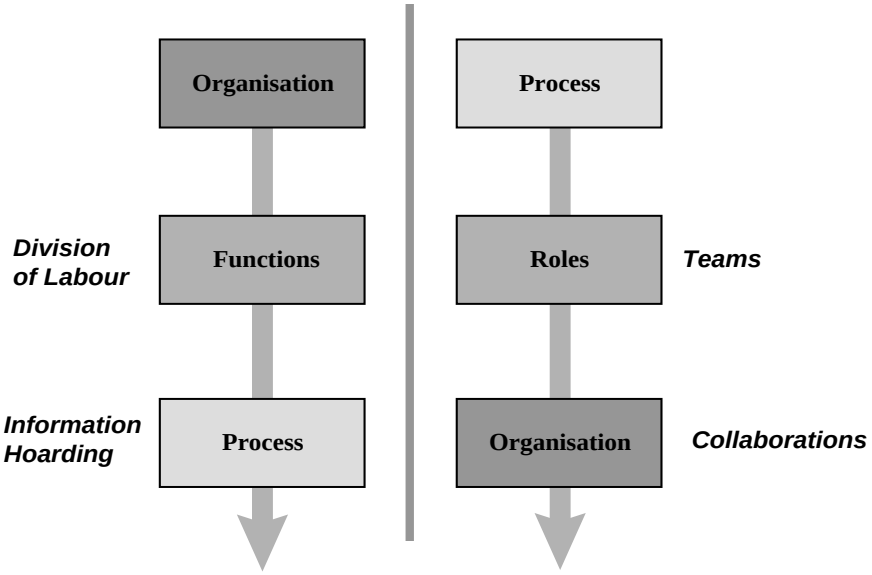

Figure 3 - Dismantling the Industrial Age
(source: ICL & Nortel)

This emphasises the move away from functional organisation towards virtual teams and processes supporting business collaboration.

Since responsibility is increasingly defined in terms of role rather than person, there is a requirement for workflow systems to maintain an audit trail of the work performer who actually undertakes a particular activity. In some cases this is complemented by a supervisory role for individual activities or the overall process which is invoked if various criteria (for example deadlines) are not met. The concept of responsibility also needs to cope with activities which are wholly automated with no human involvement (for example by IT application).

5.1.2 Process Modification

Adaptive processes are fundamental to the ongoing value of workflow; in practice adaptation can occur in several ways with different associated complexities of automation.

(i) An ongoing change to the process, introduced by the owner. An example might be the introduction of some additional checking on an authorisation task, or amendment of the value at which additional checking is undertaken. The changes may need immediate application to all existing open business cases, or may just apply to new cases.

(ii) A variation to the normal process behaviour may be pre-defined under certain conditions, for example an activity may be skipped or delegated to a subordinate role if a certain business criteria is met. This variation is defined as part of the persistent business rules applied within the process behaviour, but needs to be separately monitored (and reported) in each case.

(iii) In some cases process behaviour (i.e. the business rules) may be arbitrarily adapted or developed during operation, without any prior constraints imposed when the process was designed. This is a behaviour pattern often associated with ad-hoc workflow systems in co-operative workgroup situations, where only a skeleton process may be defined within the process definition. This is amended dynamically as process execution proceeds to add new tasks or amend responsibilities, etc.

5.1.3 Process Structure & Organisational Boundaries

One consequence of the flattening of organisational structures and increasing business integration across organisations is that process scope is extending not just across departmental boundaries but also between enterprises. A model of workflow put forward by the Japan Standards Association (JSA) Groupware Committee (1997) illustrates this industry direction by a three tier framework embracing workflow at departmental, enterprise and inter-enterprise levels.

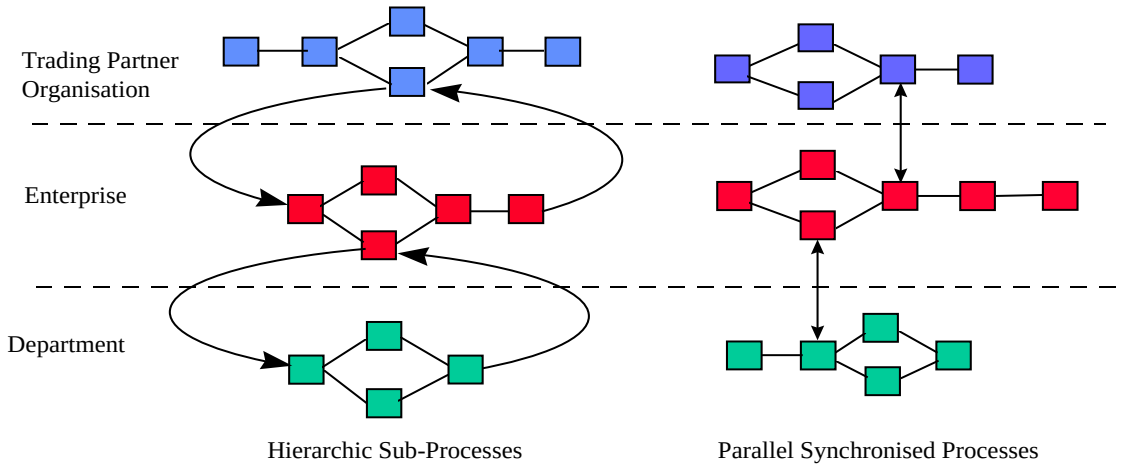

Figure 4 - Inter-organisation Business Processes

This leads to requirements to support process structures which:

1. enable a sub element of a process to be initiated in a different organisational domain (hierarchic or chained sub-processes)

2. support the periodic synchronisation of process activity between two (or more) essentially independent processes operating in different domains (parallel synchronised processes)

Such process models often impose additional constraints on automation in the areas of security between domains and conventions for object naming and organisational mapping.

5.1.4 Process Duration

This impacts a number of engineering issues, particularly the likely concurrency of active process instances and possible requirements for the support of a *dormant* process state. Most typical processes have a relatively short duration, typically from seconds to weeks. Some, which are customer-centric, may be defined in terms of a customer life cycle lasting many years. Since most workflow systems carry a significant overhead per process instance there may be a requirement is such cases to remove dormant cases to some form of secondary process data storage.

5.1.5 Activity navigation

One characteristic of all business processes is the thread of control which links together the various activities during the life of the process instance. Typically this involves conditional logic and a number of alternative routes (navigation paths) through the process. These paths generally need to support a mixture of sequential and parallel activities within a process.

This logic may be defined in quite different ways within different process definition methodologies:

Transition Based - Typically derived from Petri Net methodology, the process is represented graphically as a network of activity *nodes* and explicit *transitions* between them. *Edges* connect nodes to transitions (*input arcs*) or transitions to nodes (*output arcs*). Parallelism within a process is supported by transitions with multiple output arcs (a *split* into multiple execution threads transferring to different activities) or with multiple input arcs (a *join* of several execution threads into one). Alternative routes between activity nodes are evaluated by reference to *conditions* associated with the transitions. Although arbitrary complexity can be supported, multiple transitional expressions involving complex conditional evaluations can become cumbersome to represent in a machine processable form.

Block Structured decomposition - In this approach any single node in a model may be decomposed to a lower level of underlying process (a paradigm based upon the hierarchic sub process model). In this approach parallelism is constrained to operate only within the context of a single level of decomposition (i.e. parallel threads cannot transcend block boundaries). A product based upon this approach cannot cope with an arbitrary complexity of split and join constructs (for example an unbalanced split where one path continues beyond the context of the current block)

Activity Pre- & Post-conditions - In this approach no explicit transitions between activities are declared. The process is defined as a set of activities each having entry (*pre*-) and exit (*post*-) *conditions*; parallelism is implicit and when pre-conditions are met the activity is initiated, independently of the status of other activities within the

process. To provide sequential operation, pre-conditions may relate to the completion of a particular prior activity (and by extension to multiple prior activities, providing an "*and-join*" capability). Post-conditions may be used to control looping within an activity.

Each of the above approaches has its pros and cons and its own particular devotees. The problem for the systems integrator is that it is not easy to transfer process information between design tools and/or workflow control software based upon the different design paradigms.

5.1.6 The Organisational Model

Virtually all business processes are based around the concept of an individual's roles and responsibilities for the various activities within the process. As far as possible processes need to be isolated from the vagaries of organisational change, leading to the requirement for a (dynamic) organisational model. This can map the ongoing roles and responsibilities at process level against the current organisational entities and the current set of individuals who undertake the various roles.

The identification of an activity "performer" within a process may embrace a mixture of organisational and role information (… "the fault analyst in the European Customer Support Unit"). Organisational relationships often expressed include:

- manager of
- deputy to
- alternative to (proxy)
- role or skill profile

Responsibility models may introduce additional constraints on work performers (for example .. "not the person who authorised the original loan"), which require process history to be maintained.

Thus in many cases an organisational model will need to accompany a business process to enable its automation.

5.1.7 Security

Security is often developed separately from the business process and may have to added at automation stage by reference to a separate organisation security policy document. Many of the security requirements during automation will be related to roles, responsibilities and authority within the process.

5.2 Representing the Business Process

In order to provide automated support for a process, it must be first be captured in a machine interpretable representation. This representation must have the flexibility to structure and maintain all the process related information necessary to enable co-ordination of enactment using IT infrastructure.

The WfMC glossary introduces the term "*Process Definition"* for this representation, describing it thus:

"The automation of a business process is defined within a Process Definition, which identifies the various process activities, procedural rules and associated control data used to manage the workflow during process enactment".

The process definition may be represented by a combination of any or all of textual script, graphical notation, or formal programming notation, with many different process development tools available to manipulate such information. Typically their use follows a cycle of analysis, modelling, implementation, feedback and further analysis.

Several attempts have been made to define a standard representation of all or part of a process specification.

IDEF [5] is a series of modelling notations introduced by the US Air Force, several of which are published as FIPS by NIST. Included are methodologies for modelling business functions (IDEF0), information models (IDEF1X) (both widely used), dynamic system behaviour (IDEF2) and Process Description Capture (IDEF3).

CDIF [6] has defined a core architecture for CASE tools and data interchange bindings, based around a meta-meta-model. Foundation and Common meta-models are defined and work has been completed on data definition, data flow and data modelling. An extension to cover business process modelling is under discussion, but work is not yet mature. **UML** (unified modelling language) is a similar initiative under the auspices of the OMG, with its own modelling notation and meta-model.

None of the above currently provide a machine processable process definition as a basis for workflow automation.

PIF (Process Interchange Format & Framework) [7] has been developed by a working group drawn from a number of US and UK universities. Its underlying philosophy is that of generality over computational efficiency; this is reflected in the organisation of its entity classes which is not necessarily well suited to the performance of any specific task, such as workflow management or process simulation. It has been used for experimental translation of process related information within the research group. As with other process representations it has been found necessary to structure into a minimal core set with add-on classes. PIF is designed to be machine processable, but is not specialised to the entities and attributes required for workflow.

The **NIST PSL** (Process Specification Language) group [8] is a study group formed by NIST in April 1997, working towards a common process specification language, rather than interchange format. It has members drawn from industry, government and academia but has a particular interest in the application of process technology to manufacturing industry. There is no current specification produced by the group, but it is reviewing inputs from other industry organisations.

WPDL (Workflow Process Definition Language) [1] is specified by the WfMC and despite its name was conceived as a text-based, machine processable interchange format, rather than a definition language.

The WfMC has produced a process definition "meta-model", shown below, which attempts to capture the highest level objects and relationships which, as a minimum, must be defined to support process automation. This meta-model underpins the WPDL grammar.

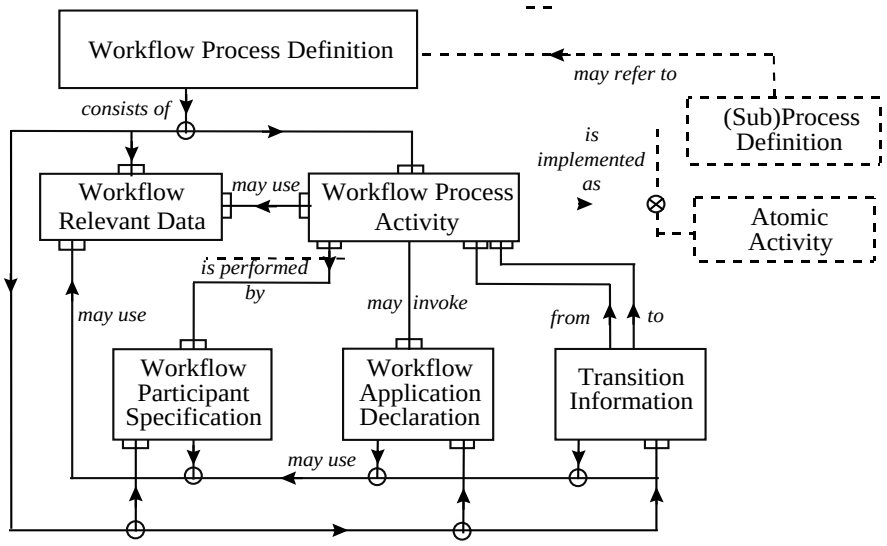

Figure 3 - The Process Definition Meta-Model

The route followed by the WfMC is to define as standard attributes the most commonly required properties of these top level objects, but to allow extensibility through an extended attribute list and library functions within the WPDL grammar.

The model and the WPDL constructs are focused specifically on workflow and provide more detailed structures for defining the workflow related aspects of a process. They do not attempt to incorporate the level of generality of other approaches such as PIF. WPDL will shortly be released in beta form and several prototype implementations have been made against interim specifications with reasonable success.

One key difficulty with all approaches remains the capture of all the potential dynamics of a business process within a single model. It is probable that automatic translation of 100% of such business processes between different products is an unreachable goal in the foreseeable future. However the "*meta-model*" provides a structure for mapping a large part of the business process into WPDL or, potentially, other interchange forms.

6. The Systems Integration Model

The WFMC is the principal organisation defining standards for workflow and is attempting to cover much of the ground discussed in earlier sections. The standardisation programme is based upon the "*Reference Model*" [1] shown below.

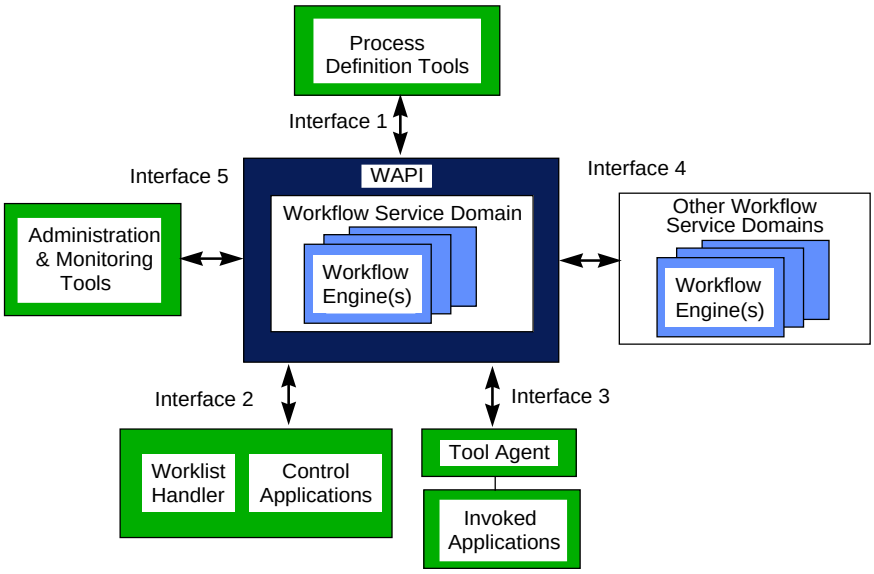

Figure 4 - The WfMC Reference Model

Whist this is an oversimplification of workflow from the integration perspective, it has proved useful within the industry as a focus for standardisation work. It concentrates on modelling a workflow service as a black box object viewed from its external interfaces, whilst ignoring the internal construction architecture (and hence a number of the integration problems).

The internal components of the "*workflow service*" are assumed to be homogeneous, and typically supplied by a single vendor's product(s). This avoids issues associated with service administration and security, which essentially lie inside the "box". Also no distinction is made between a single centralised "*engine*" and co-operating, distributed engines, which need to support shared process state data in order to support a single homogeneous service image. The model also attempts to avoid dependence on the nature of the underlying distributed infrastructure through the specification of APIs, or interchange formats, by which system components interact, which are assumed to be supportable through the infrastructure[6].

Five "interfaces" are identified within the Reference Model, realised by a combination of APIs, protocol and format conventions.

6.1 Process Definition Interchange

The purpose of this interface is to support the exchange of process definition information between BPR tools, workflow systems, and process definition repositories. The interface is based upon the meta-model described in Section 5; information exchange is supported in two ways:

1. the WPDL grammar supports the transfer of complete process models via file transfer, typically using an import/export mechanism from native product formats. The import process can check the process model for structural integrity, for

[6] One exception to this is the interoperability protocol between workflow domains, which is discussed later.

example flagging isolated activities with no transitions. The export process must flag any structures which cannot be represented in WPDL.
2. APIs are defined for reading & writing individual object & attribute data within the Process Definition. These are typically used for ad-hoc process modification or control functions, rather than bulk process transfer. There is no automatic mechanism for checking the integrity of the resultant modified process.

6.2 Client Applications Integration

This provides a standard interface for work allocation to the desktop environment, allowing desktop applications portability & re-use across different workflow environments. APIs are defined for:

1. Process & Activity Control functions, for example starting, suspending, terminating a process instance or sets of process instances
2. Worklist Handling, to allow users to log on and process (or re-assign) individual work items

6.3 Applications Invocation

This provides a single interface which may be used for two purposes:

1. To provide a common framework for the integration of software agents providing access to other industry services such as document repositories, meeting schedulers and email which use their own specific industry APIs
2. To support access to legacy applications via application specific methods (for example terminal emulation or proprietary TP protocols)

A simple API set supports Connect/Disconnect, Invoke Application, Request Status and Terminate Application

6.4 Process interoperability

This supports the remote invocation of a sub-process on a different workflow system, allowing a single business process to be implemented over two or more workflow systems.

Two variants of interchange protocol are defined:

1. MIME (Multipurpose Internet Mail Extensions) defined in RFC 1341
2. IDL bindings for use with CORBA (typically via ORB interoperability services)

This interface uses essentially the same API set as that for process initiation from client applications[7].

A missing element in the current specifications is support for synchronisation points between parallel execution threads; this is identified for future development.

The sub-process interoperability model as currently specified makes no requirement to dynamically share state data between the two interoperability domains and specifies a minimal level of prior co-ordination. (This is essentially limited to a knowledge of the called address for a particular sub-process and any related security attributes.) Thus it is more suited to “loosely coupled” distributed process enactment

[7] Since remote invocation can occur via an asynchronous interface such as e-mail some additional optimisations are provided to allow grouping of calls (or call responses) into an underlying MIME transfer.

across different organisational entities than tightly bound workflow systems within a single department or workgroup.

Some issues of detail still remain, for example which properties of sub-process operation are inherited from the superior calling process and which from the local process definition. In general it is an accepted principle that where remote process "hand-off" occurs it will not be feasible to retain all process attributes through the call and return. Details of name space usage across the two environments also remains to be fixed in detail.

6.5 Audit and monitoring

Auditing is an important requirement for many workflow systems. This "interface" comprises a specification of standard audit events and their recording format, thus enabling the integration of audit trails across different systems during workflow interoperability. The means by which audit data is accessed or retrieved on any particular system is undefined but is typically via SQL for many workflow products.

APIs are also defined to retrieve status information on current process instances or activities.

7. Ongoing Standardisation Work

The standards currently defined will make a significant contribution to workflow systems integration - provide they are adopted by product vendors. An important indicator of intent is the current OMG standards approval cycle for workflow technology [9], in which the WfMC standards are currently supported by more than 35 organisations.

Various important extensions have been identified to improve the potential of the model as a basis for integration.

Object Integration - the work with OMG has identified potential requirements for developing the architecture "internal to the workflow manager", to facilitate the integration of other complimentary OMG services such as OTS (transaction services), naming, security and versioning, etc. This approach can support closer integration between different workflow products where all use the same underlying object services architecture. There is also interest in positioning workflow within the OMG *Business Objects Framework* to identify reusable service elements which can be consolidated into a business application environment.

Security - the approach here is to specify how existing security standards should be applied in the context of workflow. The most important area is seen to be the use of authentication, integrity and confidentiality services applied to workflow interoperability, particularly between domains in different organisations.

Support for Event synchronisation - Event synchronisation represents a significant extension of the interoperability model to support transitions (and potentially associated data flow) between different, essentially independent processes, running in

different domains. Issues to be resolves include process, thread and event naming, and event management functions applied across distributed, heterogeneous products (e.g. to detect and prevent deadlock and persistent wait states).

Process Integrity and recovery - This is an area which has not been widely addressed and will take some time to mature. Recovery may require the basic process state data, shared workflow and application data and wholly application related data (for example within legacy applications). Different techniques include 2-phase commit and rollback (whose use may be impractical through asynchronous messaging infrastructure and/or long transaction times), compensating transactions, or alternative transactions. Many products rely on at least some manual recovery elements.

Internet and electronic commerce - There is considerable interest in support for inter-organisational workflow functionality carried via the Internet. Existing functionality via electronic mail will be augmented by support for more dynamic process binding (for example using traders or yellow page services). The use of XML[8] to encode process based exchanges is also under discussion.

8. The future

We shall not know for a year or two whether this standardisation programme will really contribute to the integration task. There are encouraging signs that the industry has recognised the benefit of a common architectural framework to assist with product interoperability and most products are structured in broad alignment with the reference model. In practice the number of conflicting products is bound to fall, if only through market consolidation. Interest continues in capturing and reusing automated process fragments within an applications framework architecture.

The notion that workflow will evolve into ubiquitous middleware, in the same way as, say, electronic mail, is perhaps more questionable. This requires both standardisation and a re-orientation of commercial thinking towards the value of automated processes. There is certainly every likelihood that workflow interoperability will substantially increase in inter-company trading situations. A demonstration of an automated supply chain process scenario [10], in which the overall business process was automated across 7 diverse organisational systems, attracted huge industry interest[9].

Within this style of operation it is possible to enact business processes which automatically call other organisations to implement those parts of the process which lie within their domain of responsibility , for example manufacturing, wholesaling or supply logistics. Such business interactions go far beyond the simple transfer of order data or supply notes, bringing opportunities for expressing a complete supply chain business logic in a manner which can be seamlessly automated across diverse business entities. This may well point the way towards a second generation of

[8] XML (Extended Markup Language) - a more generalised version of HTML, also derived from SGML principles

[9] Workflow Canada, Toronto, June 1996, and repeated at the Giga Workflow conference, Amsterdam, October 1996

electronic commerce based on process interoperability rather than simple electronic data interchange.

Acknowledgements
Many colleagues within the WfMC have contributed important ideas to the subject of workflow integration and interoperability. Particular thanks are due to Mike Anderson, from the TeamWare Integration Centre, who provided comments on this paper and was also responsible for the WfMC interoperability specifications.

References
[1] Workflow Management Coalition documentation:
Glossary, 1996, *The Workflow Reference Model, 1995, Workflow API Specification, 1995, Workflow Interoperability Specification, 1996, Process Definition Interchange Specification (draft), 1998.*
Details available via http://www.wfmc.org
[2] Open*framework, Creating Potential for Change*, Technical Journal 8(3), 1993
[3] Ovum Group, *Ovum Evaluates Workflow*, 1996
[4] GIGA Group, Conference Proceedings, Business Process & Workflow, London, 22-24 October, 97
[5] Koulopoulos, Thomas, *The Workflow Imperative*,1995, Van Nostrand Reinhold, ISBN 0-442-01975-090000
[6] IDEF (Integrated Computer Aided Manufacturing Definition), details of IDEF0, IDEF1X, IDEF3, plus work in progress available via the IDEF home page http://www.idef.com
[7] CDIF (CASE Data Interchange Format) specifications available via http://www.cdif.org
[7] PIF (Process Interchange Format), details available via PIF home page *http://soa.cba.hawaii.edu/pif/*
[8] NIST Process Specification Language (PSL) project documents:
Proceedings of the First PSL Roundtable, NISTIR 6081, National Institute of Standards and Technology, Gaithersburg, MD (1997).
Unified Process Specification Language: Requirements for Modelling Process, NISTIR 5910, National Institute of Standards and Technology, Gaithersburg, MD (1996).
Both available via NIST PSL home page *http://www.mel.nist.gov/psl/*
[9] Object Management Group, Framingham, MA 01701-4568, RFP for Workflow Technology, 1997 and associated proposals, details via OMG
http://www.omg.org/library/schedule/Workflow_RFP.htm
[10]Anderson, M.J., *Workflow Engine Interoperability - What's in it for Users*, Document World (May/June 1997)