

---

## Integrating e-government public transactional services into public authority workflows

---

Costas Vassilakis\*, George Lepouras,  
Stathis Rouvas and Panagiotis Georgiadis

Dept. of Informatics and Telecommunications, E-Gov Lab,  
University of Athens, Panepistimiopolis, 15784, Athens, Greece

E-mail: costas@e-gov.gr E-mail: gl@e-gov.gr

E-mail: rouvas@e-gov.gr E-mail: p.georgiadis@e-gov.gr

\*Corresponding author

**Abstract:** Documents submitted by citizens through electronic services, deployed in the context of e-government, must usually undergo processing by some organisational information system in order to complete the citizens' requests and for the reply to be returned to the citizen. The integration, however, of the e-service delivery platform and the organisational information system is often hindered for a number of reasons, including security considerations, platform diversity or idiosyncrasies of legacy information systems. In this paper, we present a generic method for providing seamless communication between the two platforms, enabling the full integration of documents submitted through electronic services into the organisational workflow, thus leveraging the quality of services offered to citizens and facilitating e-service development and operation.

**Keywords:** electronic services; e-government; integration; communication; security; scheduling.

**Reference** to this paper should be made as follows: Vassilakis, C., Lepouras, G., Rouvas, S. and Georgiadis, P. (2004) 'Integrating e-government public transactional services into public authority workflows', *Electronic Government*, Vol. 1, No. 1, pp.49–60.

**Biographical notes:** Dr Costas Vassilakis is currently a visiting Associate Professor in the department of Computer Science and Technology of the University of Peloponnese. He holds a BSc in Informatics (1990) from the Department of Informatics, University of Athens and a PhD from the same department in 1995. Dr. Costas Vassilakis has participated in several European and national projects and has published more than 30 scientific papers in international journals and conferences. He has been a consultant for the General Secretariat for Information Systems of the Ministry of Finance, Greece.

Dr George Lepouras is currently a visiting Assistant Professor in the department of Computer Science and Technology of the University of Peloponnese. He holds a BSc in Mathematics (1991) from University of Athens, an MSc in Information Technology (1992) from University of Strathclyde, Scotland and a PhD in Human-Computer Interaction (2000) from University of Athens. Dr. Lepouras has participated in several European and national RTD projects and has published more than 30 scientific papers in international journals and conferences. He has also been a consultant to the General Secretariat for Information Systems of the Greek Ministry of Finance.

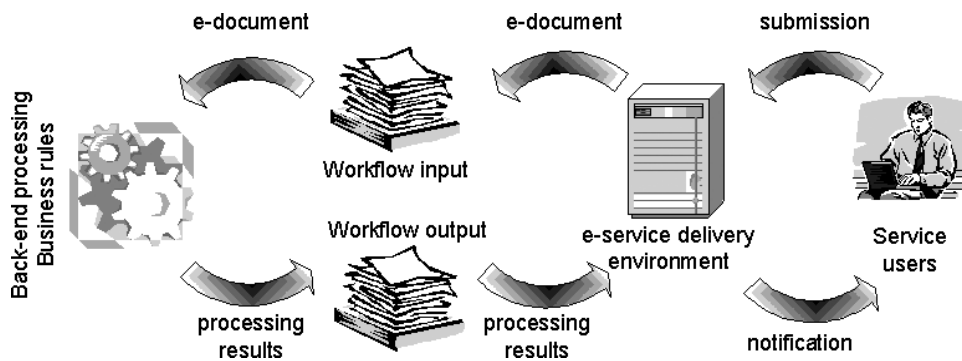
Mr Stathis Rouvas was born in Athens, in 1968. He holds a degree in Civil Engineering from the National Technical University of Athens, Greece and is now a PhD candidate in the Department of Informatics and Telecommunications, University of Athens. Mr Stathis Rouvas has participated in European and national RTD projects, and is also a consultant to the General Secretariat for Information Systems of the Ministry of Finance. He has published more than 10 papers in national and international journals and conferences.

Dr Panagiotis Georgiadis is an Associate Professor in the Department of Informatics and Telecommunications of the University of Athens, Secretary General for e-government and head of the e-gov lab. He holds a BSc in Physics, an MSc and PhD in Computer Science and has been a regular member of the Senate of University of Athens, Director of the Computer Systems & Applications Division of the Department of Informatics, Secretary General for Information Systems by the Greek Ministry of Finance (1997–2002). He has authored more than 50 scientific articles in international journals and conferences, and contributed in national and European research projects.

## 1 Introduction

According to the European Commission [1] “*transaction services, such as electronic forms, are perceived as the future of electronic government*”. Transaction services, in general, allow users to submit electronic documents containing the data they want to communicate to the public authority (PA) delivering the service. These data are then transferred to the PA’s organisational information system, where they are processed according to the workflow defined by the PA’s business rules. When the processing is complete, a completion notification should be returned to the citizen who initially submitted the document, containing information about the result of the processing and/or for further actions that must be performed by the citizen. The full processing cycle in the context of transaction services is illustrated in Figure 1.

**Figure 1** Full processing cycle for transaction services



The procedure for forwarding electronic documents submitted by a citizen to the organisational back-end and for relaying the results of the processing from the back-end to the citizen through the service delivery environment is not always fully automated, leading to delays in the processing cycle and a need for human intervention. The main reasons for not automating this procedure are the following:

- *Security*: an on-line connection channel between a publicly accessible service delivery environment and the organisational information system is a potential security hazard, since it may be exploited by attackers to penetrate the back-end systems. Often, proprietary and undocumented protocols are used to implement the communication, thus inhibiting the enforcement of strict firewalling rules (address/port-based and/or content-based [2]) that would alleviate the problem.
- *Platform diversity*: in most cases the electronic services delivery platform is developed and managed separately from the back-end IT system, thus, different software and hardware products are used in each environment. This may lead to various compatibility problems, such as information modelling mismatch, incompatible national character set encoding, different settings in number precision etc.
- *Inflexibility of legacy systems*: most back-end systems are designed to *deliver services only*, i.e. wait for incoming requests and execute the designated operation. In order to support the complete processing cycle for transaction services, these systems must act as *clients* as well, arranging for 'pushing' the results obtained from processing the submitted e-documents to the service delivery environment, which will produce the appropriate notifications to users.
- *Task scheduling complexity*: some e-documents must be immediately processed by the back-end systems, since the user is waiting for an on-line response; in other cases, document processing may require human intervention or it would be preferable to be deferred to some off-peak time frame. Legacy systems often support only one type of task scheduling (immediate or batch), impeding the realisation of optimal scheduling policies.
- *Change management*: when changes in the delivered electronic services occur, usually, the service delivery environment is readily updated whereas the back-end IT system is adapted at a slower pace. This lack of cohesion could introduce problems in an automated communication scheme.

In this paper, we present a communication scheme between electronic services delivery platforms and back-end IT systems that fully supports the processing life cycle of transaction services, allowing documents that have been submitted electronically to be integrated into the organisational work flow. The proposed communication scheme tackles all issues that hinder the establishment of direct communication between the two platforms, thus facilitating fully automated processing, shorter turnaround times and removing the need for human intervention.

The rest of the paper is structured as follows: Section 2 presents related work in this area. Section 3 outlines the architecture for the proposed communication scheme, Section 4 evaluates the approach by illustrating how the issues discussed above are tackled and, finally, conclusions are drawn in Section 5.

## 2 Related work

So far, work targeted to the provision of a generic framework for the integration of web services and organisational workflows has not been reported. Organisational needs have been addressed on a case-by-case basis by system implementers and integrators, who employ a number of paradigms facilitating communication between (possibly diverse) platforms.

One of the first approaches used for implementing communication between different platforms has been the *remote procedure call* paradigm (RPC [3]), which allows programs running on a specific machine to execute code fragments on a different system, in a fashion similar to local procedure calls. The RPC paradigm, however, does not incorporate strong security mechanisms and TCP/IP port assignment is not fixed, hindering the use of strict firewalling rules. Moreover, resilience against transient failures is minimal, no scheduling facilities are provided and both systems must be 'in sync'. CORBA [4,5] is a more complete approach providing a middleware layer with features addressing platform heterogeneity, reliability and failure resilience; however, CORBA-based systems do not generally support scheduling policies and change management, while most versions tend to be more complex and resource-consuming than is actually needed in an electronic service delivery environment. Additionally, CORBA-based systems are difficult to protect using firewalls and a compatible ORB is required on both ends.

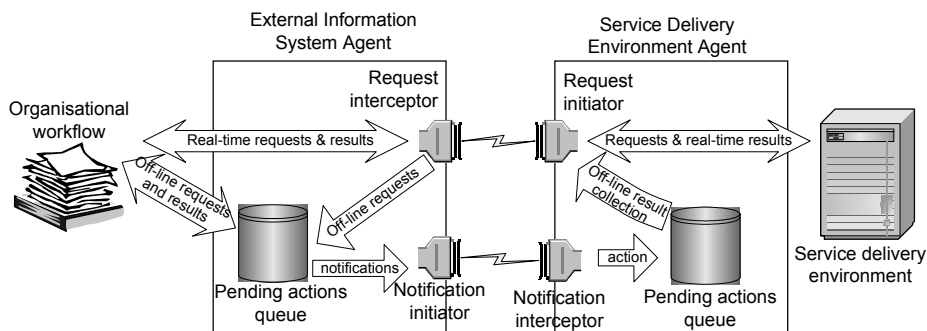
Recent developments for platform communication include the XML language [6], the SOAP protocol [7] and the XML-RPC paradigm [8]. The XML language is the emerging standard in platform information exchange, since it allows for incorporation of structure and semantics in the exchanged messages, but this standard only defines the syntactic language rules and does not deal with issues such as security, scheduling or, more generally message handling. SOAP is an XML-based messaging protocol, defining a set of rules for structuring messages that can be used for simple one-way messaging and for performing RPC-style request-response dialogues. The SOAP approach involves a listener but security provisions are minimal, whereas no scheduling alternatives are provided (requests are always executed immediately) and no facilities are provided for handling change management. Finally XML-RPC is a variant of the RPC paradigm, where request parameters and responses are encoded as XML messages instead of using the marshalling and unmarshalling procedures of standard RPC. XML-RPC may also be tunnelled through the HTTP protocol. These enhancements, however, do not address any of the aforementioned issues regarding the use of the RPC mechanism in the context of electronic services.

## 3 The communication scheme

The proposed communication scheme encompasses two software modules, namely the *Service Delivery Environment Agent* and the *External Information System Agent*. The *Service Delivery Environment Agent* is an autonomous software module running on the platform hosting the e-service delivery environment. This software module enables the submission of *requests* to the PA's back-end system (or, more generally, to any system external to the service delivery platform) and the retrieval of the respective results.

The External Information System Agent is again an autonomous software module running in the environment of the PA's IT system (or in the environment of any system that the service delivery environment needs to communicate with) and arranges for the interception of requests originating from the Service Delivery Environment Agents and placing the e-documents contained in these requests in the input queue of the organisational workflow. Moreover, the External Information System Agent undertakes the responsibility for collecting the responses from the output queue of the organisational workflow and delivering them to the service delivery environment for further actions to be taken. Two *pending action queues* provide facilities for deferred scheduling of non time-critical requests, and for providing resilience against temporary communication failures. The overall architecture of the proposed communication scheme is illustrated in Figure 2. The operation of the communication scheme is described in the following paragraphs.

**Figure 2** Overall communication scheme architecture



### 3.1 Communication scheme functionality

The proposed communication scheme allows for services running within the service delivery environment to submit requests for e-document storage and retrieval from the back-end system. In order to submit a request, the electronic service should use an appropriate API, which is implemented in a library that must be loaded by the electronic service. Currently a Java [9] implementation for this library is provided and implementations in the PHP [10] and ColdFusion [11] scripting languages are underway.

Each request is characterised either as *real time*, if it must be instantaneously processed and the results immediately returned [12], or as *off-line*, if the processing of the request may be deferred. An example of a real-time request is the event that a service user requests to view a submitted document; this request should be executed immediately, since the user is waiting at his client device for the document to appear. If completion of the request is not immediately possible, an appropriate error should be reported. On the contrary, when the user has filled in and submitted a new e-document, the organisational policy may allow a two-day time window for the e-document to be processed in the organisational back-end and the reply to be returned; in this case, the execution of this request within the organisational back-end may be deferred. The characterisation of a request as *real-time* or off-line is dependent on the request semantics and organisational policies.

The *request initiator* module of the Service Delivery Environment Agent arranges the forwarding of these requests to the External Information System Agent, where the *request interceptor* module receives it. This communication is performed using either RMI [13] or XML [6] messages. In both cases, transferred messages may be encrypted by employing SSL or TLS [14].

Upon receiving an incoming request, request interceptor inspects the associated real time characterisation. Requests tagged as *real-time* are immediately forwarded to the organisational workflow for processing, the results are collected and returned to the service that submitted the original request. This mode of operation is equivalent to the on-line services offered by databases, transaction monitors etc.

The actual method of communicating with the organisational workflow for the purposes of submitting or extracting electronic documents is heavily dependent on the implementation of the workflow. In some cases, database tables must be appropriately queried, to retrieve documents, or populated, in order to store documents; other implementations require the execution of a custom program (usually provided with the workflow environment) that will read or modify the workflow e-document repository. Finally, certain workflow engines provide programming interfaces (APIs) through which programs may interact with their electronic document repository. The *request interceptor* handles these cases through appropriate configurations, as detailed later in this section.

When an incoming request characterised as *off-line* is received, the request interceptor places it in the *pending actions queue*, a system repository residing in the environment of the organisational information system, and notifies the requesting application that its request has been successfully received and stored. A dedicated procedure, the *pending actions queue dispatcher*, periodically scans the pending actions queue to determine if there exist jobs that can be executed. The decision as to whether a specific job in the pending actions queue will be scheduled for execution depends on the following parameters:

- *Current system load*: system administrators may configure the pending actions queue dispatcher so as not to spawn new tasks when the system load has exceeded a certain limit. This facility is provided to enable system administrators to favour the execution of real-time tasks (which are more time critical) and to guard the system against overloading.
- *Organisational policies for batch job execution*: the organisational policy may specify certain time slots for batch job execution, such as weekends, night hours or off-peak periods. In a more refined environment, different policies may apply to different tasks; for instance a request to simply store a tax return document may be allowed to run on any day of the week, but requests to process such documents and compute the tax due may only be allowed to run during weekends. The pending actions queue dispatcher can be configured to take into account such policies and appropriately schedule off-line requests.

Once the pending actions queue dispatcher has determined that a specific request in the pending actions queue can be scheduled within the organisational workflow, the suitable method is employed for storing or retrieving the relevant document, similar to the case of real-time requests, described above. An issue raised in this case is that when the processing of a request has been completed, the electronic service that originally submitted the request may not be active, thus being unable to collect the response.

For example, if a citizen submits a tax return form on Monday and the computation of the tax due is performed at a weekend, it is unrealistic to assume that the citizen will remain on-line for this period to collect the reply. Therefore, replies to off-line requests should be spooled and returned to the service delivery environment in an asynchronous manner. Delivery of replies to off-line requests is facilitated by employing the *pending actions queue* in cooperation with a signalling mechanism. More specifically, when the pending actions queue dispatcher forwards an electronic document to the organisational workflow, it arranges to collect the reply and signal to the service delivery environment that the request has been processed and the reply is ready. The reply is stored, with appropriate tagging, in the pending actions queue and a *notification event* is posted to the service delivery environment to signify that the result is ready.

A *notification event* contains all information that is required by the service delivery environment to determine which result has been prepared and its mapping to the initial request; special provisions have been made for the case where a single 'prepared result' contains the replies to multiple requests, which have been packaged into a single container for practical or optimisation purposes. Two software modules undertake the delivery of the notification event to the service delivery environment, namely the *notification initiator*, running on the installation hosting the organisational information system, and the *notification interceptor*, running on the premises of the service delivery environment. The *notification initiator* module accepts requests for posting notifications to the service delivery environment, verifies their correctness and arranges for their delivery to the *notification interceptor*. Upon receipt of a notification event, the *notification interceptor* places a suitable entry in its local pending actions queue (different from the pending actions queue residing in the premises of the organisational information system). A local *pending actions queue dispatcher* examines this queue and processes its entries in order, so as to fetch the results of off-line requests in the service delivery environment, thus completing their processing cycle. Processing of an entry of this pending actions queue usually maps to the execution of a real-time request, specifically crafted for the purpose of collecting the specific type of results.

This mode of operation is the counterpart of batch process execution offered by most legacy systems. For requests serviced in this mode, there is no need for the organisational workflow to provide fully automated procedures for their execution, since the results are not expected immediately in the context of the delivered service.

The pending action queues in the premises of the service delivery environment and the external information system are also used to provide resilience against transient errors. More specifically, when an off-line request cannot be forwarded from the service delivery environment to the external information system due to a communications failure, the request initiator stores the request in the pending actions queue and its transmission is retried after a period of time. Real-time requests, on the other hand, are not stored in the pending actions queue, since their semantics dictate that immediate execution is called for. On the side of the external information system, notification events that cannot be posted due to communications failures are similarly stored in the local pending actions queue and their transmission is retried after a time interval.

### 3.2 Configuring the communication modules

In order to provide a generic and flexible communication framework fulfilling all requirements for information exchange in a service delivery environment, each module of

the *service delivery environment agent* and the *external information system agent* is fully configurable through a set of parameters:

- 1 The request initiator may be configured to selectively forward some request types to a specific back-end system (through the associated *external information system agent*), while some other request types may be forwarded to another back-end system. For instance, in an environment delivering taxation services, requests for storing documents containing VAT statements may be forwarded to different back-end system than requests for storing documents containing tax return forms. Moreover, it is possible for system administrators to define *alternative execution methods* for a specific request. In such a case, the request initiator module will sequentially try the defined methods, until an attempted method succeeds. This feature is provided to facilitate usage of *failover systems*, i.e. systems that organisations install to undertake the provision of specific functionalities when the ‘normal’ service provider is not functioning or is unreachable. For example, a municipality offering electronic services may install two systems, *S1* and *S2*, hosting the citizen registry, and configure the request initiator to forward registry-related requests *first* to the system *S1* and, if this attempt fails, to system *S2*. Under this configuration, normal operational requests are served from system *S1* (the *primary system*); if system *S1* is not operational (e.g. due to malfunction, maintenance or communications failure), and thus the request initiator will fail in its attempt to contact it, requests will be forwarded to system *S2*. It should be noted that the request initiator does not undertake any responsibility for synchronising the information repositories of multiple systems providing the same service, in the presence of update requests; information repository synchronisation should be provided by the systems themselves. Finally, the ability to define alternative execution methods *without any order* is provided; in such a case, the request initiator will forward the request to a randomly selected service provider from within the available pool. This configuration set-up may be used for implementing load balancing.
- 2 The request interceptor may be configured regarding the details of how each incoming request is integrated in the workflow of the organisational information system and how the relevant results may be retrieved. Since workflow implementations radically differ in this aspect, the request interceptor can be configured to support different methods of interacting with the workflow engine:
  - *Database access*: when this option is used, the notification interceptor will store incoming electronic documents in designated database tables or will query database tables to produce electronic documents as results. An option for executing stored procedures is provided as well.
  - *External program invocation*: this option provides the facility for executing an operating system level command, possibly complemented with command line parameters. In this case, the electronic document is fed as the input to the invoked program and the program output is collected as a reply to a request. Wrappers are also provided for external programs needing to read the e-document or store it to an operating system file.



- *Procedure invocation*: for workflow engines providing a concrete API for querying or populating the electronic document repository, the request interceptor provides the facility for invoking a procedure, written in the appropriate programming language, which will arrange to appropriately interact with the information repository to store or retrieve electronic documents. This feature may only be exploited in environments supporting dynamic code loading, such as the Java platform and operating systems with provisions for dynamic library loading.
- *File storage*: if none of the above methods is applicable to the workflow engine at hand, the request interceptor may store the incoming request to an operating system file, for further processing in semi-automatic or manual manner.

The request interceptor may be additionally configured to support *alternative execution methods* for each type of incoming request. Similar to the case of the notification initiator, alternative execution methods are tried sequentially until one of them succeeds. This provision allows the exploitation of multiple access paths to the workflow repositories that may be provided.

- 3 The notification interceptor may be configured to designate which action(s) should be scheduled for execution, as a response to each kind of notification event received. Action designation is performed by placing relevant entries in the *service delivery environment* pending actions queue, which is then read by the pending actions queue dispatcher.
- 4 For the pending actions queues (both for the one hosted in the service delivery environment and for the one hosted in the external information system), the system administrators may configure the period that entries are examined, the clock hours during which each type of entry may be processed and the system load over which execution of pending tasks should be precluded.

#### 4 Assessment of the proposed scheme

The proposed scheme has been installed and tested in the context of two public authorities offering electronic services. A mixture of electronic services with diverse requirements was selected, in order to verify that the scheme meets the needs posed by organisations. In more detail, the selected services included the following:

- *Submission of informational documents*: i.e. documents that needed to be submitted by citizens in order to notify the PA authority of certain facts. These documents have no automated or semi-automated process associated with them and produce no reply to the citizen. PA inspectors use them for crosschecking purposes in samples of the population.

In these cases, three generic services were required: creation of a new document (blank, except for the user data which were filled in), submission of a document and viewing of a document. In one case, document editing was also required. All services were coded as synchronous requests, since users should be able to view the documents they submitted immediately. If this requirement was relaxed, the storage request could be coded as an asynchronous one.

- *Submission of instantly processed documents:* this case includes documents that are submitted by citizens and are instantly processed to produce a reply for the submitting citizen. Submitted documents are synchronously forwarded to the back-end where they are processed and the reply is returned as the response to the 'submit document' request. The collected reply is presented to the citizen. In all cases, the back-end procedure that processed the document also arranged for its storage, thus saving one call from the service delivery environment to the organisational back-end. This optimisation step requires that the system responsible for processing the document can access the repository into which submitted documents are persistently stored; if this is not the case, two separate requests (for processing and storage) should be coded in the service delivery environment.
- *Submission of batchly processed documents:* in this case, documents submitted by citizens are not immediately stored or processed; rather they are spooled and processed in a batch manner, either periodically or when a certain document volume has been amassed. For these cases, document processing requests have been coded as asynchronous calls, whereas document storage tasks have been coded either as synchronous or asynchronous requests, depending on the viewing policy. For the cases where a submitted document should be immediately retrievable by the citizen (for viewing/confirmation purposes), storage requests were synchronous (and consequently separated from processing requests). For the cases where submitted documents should be retrievable only after the processing has been completed, only a single request was coded and the back-end process arranged for the storage of the citizen's document in the PA's document repository. When the processing stage is complete, a notification is sent to the service delivery environment, initiating a synchronous procedure that collects the replies batch, which is then processed to produce suitable e-mail messages to the citizens concerned.
- *Submission of manually processed documents:* in this case, the PA's staff manually process submitted documents; the processing stage possibly involves some extra-systemic procedure, such as an on-site inspection, an interview etc. For these cases, the back-end procedure typically involves storing the document in the persistent storage and appropriately notifying (via an e-mail message or by raising a database flag subsequently inspected by the organisational applications) the employee in charge of this work. When the process has been completed, the PA worker enters the relevant data to the organisational application, via a procedure that has been enhanced to produce a notification to the service delivery environment. As a response to the notification, the service delivery environment spawns a process for collecting the relevant portions of the data, which are communicated to the citizen.

IT staff in both PAs found the communication scheme to be 'manageable' and 'easy to maintain in cases of change'. The separation of service code from the code implementing communication with the back-end system, in particular, was very positively commented on, since it provides 'a concrete task separation'. Debugging has been found to be somewhat tedious, since, in order to pinpoint an error, data from the communication scheme queues and the error logs have to be extracted and combined; however, this has been characterised as a 'typical problem in electronic services'. A dislike for XML configuration files has also been recorded, particularly from staff not familiar with the XML language.

Regarding the issues outlined in Section 1, the proposed scheme tackles them thus:

- *Security*: all communication between the service delivery platform and the organisational information system is performed:
  - 1 between the request initiator and the request interceptor
  - 2 between the notification initiator and the notification interceptor.

These software entity pairs communicate in known TCP/IP ports and exchange messages of well-documented types (XML or RMI), thus strict firewalling policies may be enforced (both address/port-based and content-based), augmenting the overall protection level of the organisational back-end system.

- *Platform diversity*: divergences owing to differences in hardware and software environments or related to different information schemata are handled by mappings encapsulated in the request initiator and the request interceptor. Messages are exchanged in XML or RMI, which provide rich data types, allowing the transfer of any type of data. In particular, XML messages may contain descriptions of the required conversions, thus facilitating dynamic mappings.
- *Legacy systems act only as servers*: the proposed communication scheme always assigns the role of the server to the back-end system and the role of the client to the service delivery environment, thus being in tune with the functioning of the back-end systems. Bi-directional information exchange is handled through notifications, which are undertaken by the proposed communication scheme, not affecting the organisational information system.
- *Task scheduling complexity*: the proposed communication scheme includes configurable scheduling engines, thus removing the need for schedulers provided by the operating system or custom written by the IT staff. The schedulers built in the communication scheme support both real-time transactions and deferred (off-line) processing.
- *Change management*: the proposed communication scheme supports flexible mappings between the front-end information schema and the workflow information schema (within the configuration of the request interceptor), facilitating change management when the two interconnected systems are separately maintained. When a change takes effect in the service delivery environment, and the respective portion of the back-end system has not yet been updated, it is possible for the administrators to configure the request initiator so that requests affected by this change are *spooled* to a local repository (e.g. a database or an operating system file), rather than forwarded to the back-end system. Spooled requests may be transferred to the back-end system when the changes in question have been incorporated into it.

## 5 Conclusions

In this paper we have presented a generic communication scheme for integrating an electronic services delivery environment within an organisational workflow. The proposed scheme screens service delivery platforms from legacy system idiosyncrasies, supports both immediate and deferred task execution, provides for flexible task

scheduling and allows for enforcement of strict firewalling rules. The proposed scheme has already been evaluated, in terms of functionality and performance, in the context of one installation hosting five simple services needing to interact with the organisational workflow. Stress tests will also be performed, in order to assess the scalability of the platform. Tight integration with tailorability providers (e.g. [15]) will also be addressed.

## References and Notes

- 1 European Commission (2000) *Public Sector Information: a Key Resource for Europe*, Green Paper on Public Sector Information in the Information Society, Available from: <ftp.echo.lu/pub/info2000/publicsector/gppublicen.doc>
- 2 Northcutt, S., Zeltser, L., Winters, S., Fredrick, K. and Ritchey, R. (2002) *Inside Network Perimeter Security: the Definitive Guide to Firewalls, Virtual Private Networks (VPNs), Routers, and Intrusion Detection Systems*, New Riders Publishing, ISBN: 0735712328.
- 3 Sun Microsystems (2002) *ONC+ Developer's Guide*, Available from: <http://docs.sun.com/db?q=RPC&p=/doc/805-7224>
- 4 Object Management Group (2002) *CORBA Basics*, Available from: <http://www.omg.org/gettingstarted/corbafaq.htm>
- 5 Bolton, F. (2001) *Pure CORBA*, SAMS Publications, ISBN: 0672318121.
- 6 World Wide Web Consortium (2001) *The XML Specification*, Available from: <http://www.w3.org/xml>
- 7 Newcomer, E. (2002) *Understanding Web Services: XML, WSDL, SOAP, and UDDI*, Addison Wesley Professional, ISBN: 0201750813.
- 8 Laurent, S., Dumbill, E. and Johnston, J. (2001) *Programming Web Services with XML-RPC*, O'Reilly & Associates, ISBN: 0596001193.
- 9 Arnold, K., Gosling, J. and Holmes, D. (2000) *The Java Programming Language*, 3rd ed., Addison-Wesley Publishing Company, ISBN: 0201704331.
- 10 Lerdorf, R. and Tatroe, K. (2002) *Programming PHP*, O'Reilly & Associates, ISBN: 1565926102.
- 11 Hewitt, E. (2001) *Core ColdFusion 5*, Prentice Hall, ISBN: 0130660612.
- 12 In this case, the related procedure in the organisational workflow must be fully automated.
- 13 Sun Microsystems (1999) *Java™ Remote Method Invocation*, Available from: <http://java.sun.com/j2se/1.3/docs/guide/rmi/index.html>
- 14 Thomas, S.A. (2000) *SSL & TLS Essentials: Securing the Web*, John Wiley & Sons, ISBN: 0471383546.
- 15 Loverdos, C., Saidis, K., Sotiropoulou, A. and Theotokis, D. (2002) *Pluggable Services for Tailorable E-content Delivery*, OOIS, pp.6–18.