

Delivering New Electronic Services in the Public Sector

The Electronic Government Interoperability Framework and the
UK Government Gateway

Paul Spencer

Abstract

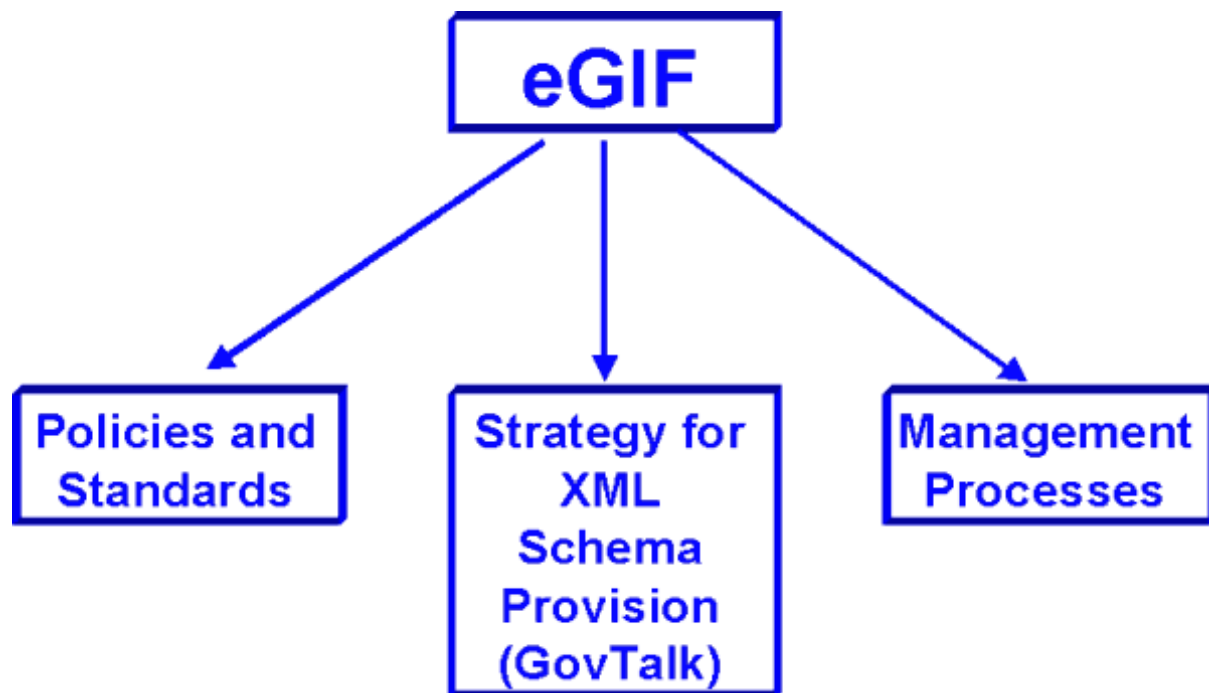
During 2000, the UK Government issued the e-government interoperability framework (e-gif) and started work on defining standards and developing systems to promote interworking between individuals, companies and national and local Government. This is all in support of the Government statement that all services will be available electronically by 2005. This paper introduces the e-gif and the work on developing XML schema guidelines and other supporting documents. It then discusses the Government Gateway - the interface into Government. This part of the paper talks about the involvement of the various commercial and Government organizations in the project, the design decisions required for such a system, the XML standards used (including the benefits and hazards of working with W3C working drafts such as XML Schema) and the final decisions taken regarding schemas and interfaces.

In March 1999, the UK Government published its 'Modernising Government' white paper (<http://www.citu.gov.uk/moderngov/whitepaper/4310.htm>) in which it announced 'a new target of making all dealings with Government available electronically by 2008'. In 1999, 2008 was so far away to those of us coming from the private sector that there was obviously no need to start yet. So, on being appointed e-Envoy, one of Alex Allen's first moves was to move the date forward to 2005. Now we were talking!

Since that date, a physical infrastructure has been designed and the first phase developed, work has started on defining new process models and an information architecture and thousands of people are working to deliver the vision. In this paper, we will look at three parts of this UK Online initiative - the Electronic Government

Interoperability Framework or e-GIF (<http://www.govtalk.gov.uk/egif/home.html>), the work being done on developing XML schemas and the Government Gateway that forms part of the physical infrastructure. At the time of writing, the e-GIF is at version 1. This paper is based on a draft of version 2. By the time the paper is presented, version 2 will have been released, so the content of the presentation may differ from this written version.

The e-GIF itself is in three sections:



The first defines the policies and standards to be used in dealings between the public and private sectors and within the public sector. The second is described as a strategy for schema provision, although it covers far more than this. The third is the mechanism for maintaining the e-GIF itself, and I will not describe that here.

The interactions covered by the e-GIF can broadly be described in three categories: citizen to government (C2G), business to government (B2G) and intra-government (G2G). In this case, government includes the full public sector including local government and the National Health Service.

1. Policies

The policies are based around use of internet standards and the Internet itself as the

means of service delivery. Given the ubiquitous nature of the Internet, this is hardly something that requires detailed justification. The other main policies are the adoption of XML as the key mechanism for data tagging and the adoption of the browser as the key interface.

The use of XML required some justification when I was promoting it to the Government in 1998, but now that decision has been well vindicated. However, in my view, there is a possibility that the e-GIF goes too far by mandating it for all transactions. There are many large inter-departmental data transfers already using EDI formats, and I question the benefit of changing these. While XML is perfect for the many-to-one and many-to-many transactions typical of the private to public sector interface, other formats could have benefits for these one-to-one transactions. This is an area likely to be corrected in version 2.

The final policy discussed here is the use of the browser as the key interface. Version 1 of the e-GIF says that all information must be available through a browser, although the definition of the browser is sufficiently broad to include WAP devices and digital TV. A browser interface should be available for document submissions to Government but it is still possible to support other interfaces. For example, when sending payroll information at the end of the tax year, small companies might like to complete a web form, while larger companies and payroll bureaux would prefer to just press a 'submit' button in their payroll system. But what about filing of company accounts and tax data, where a complete set of accounts is required? Presumably version 2 will allow Companies House and the Inland Revenue to leave it to the private sector to provide suitable enhancements to existing accounting systems rather than providing a web form which will be time-consuming to complete and could lead to transcription errors.

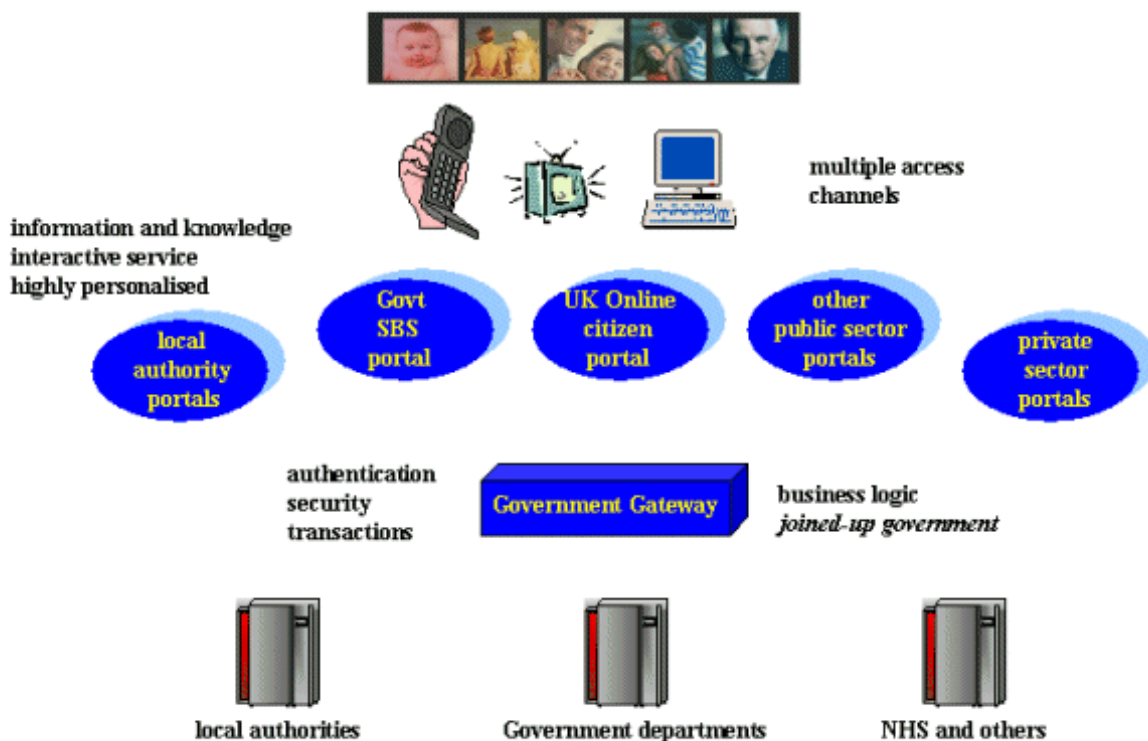
2. Standards

The e-GIF specifies a set of standards that is driven by the need for interoperability, openness, market support and scalability. Ultimately, the aim is to make support of the e-GIF simple and cheap for implementers. These implementers are not only those in the public sector, but also independent software vendors such as the manufacturers of accounting systems and farm management systems.

Three main areas are covered by standards drawn from outside the e-GIF. These cover interconnection, data integration and information access. In practice, the descriptions here are just providing more detail on the information provided by the policies. For example, the policies mandate alignment with the Internet, while the standards define IPv4 as the networking standard, while stating an intention to move to IPv6 at a future time. The standards also indicate the supported document formats. Again, these are very much those you would expect in an Internet environment, with support for HTML, JPEG, GIF and PDF amongst others.

3. UK Online Physical Infrastructure

As well as setting standards, the UK Government is installing a physical infrastructure for C2G and B2G transactions. This comprises a citizen portal, a small business portal and the Government Gateway. Of these, the citizen portal and Government Gateway are already running, albeit in early phases. The basic infrastructure looks like this:



Note that private sector web sites and portals can also access the Gateway. This means that existing services such as ihavemoved.com can add government

departments to the list of bodies that they can notify of a change of address. The advantage for operators of these sites is that there is a common format of XML message that will be used throughout the public sector.

Looking at the infrastructure from the top, we see the support of multiple access channels. These will be able to access both public and private sector portals that will hold initial conversations to gather information before submitting that information to the Government Gateway. When doing this, they might access other information in the same way that many online services find your full postal address from a house number and postcode. The Gateway looks at information in the message envelope to authenticate the sender, check that he or she is authorised to perform the transaction requested and route the message to one or more destinations. In some cases, these destinations might require different information from the document (and not be allowed under UK data protection laws to receive information in the message that is only intended for other recipients), so the Gateway has a filtering facility. Once the message reaches its destination, the response is routed back. We will look more at that process in a moment.

4. Security Requirements

The various transactions that use this infrastructure have differing security requirements. For example, no security is required for someone to request information on the vaccinations required for their vacation destination. However filing a VAT return requires a higher level of security. Currently, three levels of security are provided - none, user name and password, and digital certificates. The intention is to encourage the use of digital certificates as a way of encouraging their wider use in the private sector. This will happen first for companies and later for individuals.

The Government is not planning to issue certificates itself - it is working with the private sector in this area. It is simple for a company to get a certificate from its local Chamber of Commerce and then starting using it to communicate with Government.

Documents submitted through the Gateway can be signed with a digital certificate using the W3C specification for signing XML documents. Normally, the body of the message is signed as well as parts of the envelope to guarantee that the message that is received, processed and stored for audit purposes is the same as that

originally sent.

As well as sending documents on your own behalf, it is possible to send on behalf of someone else by registering as an intermediary (or agent). Once the relationship between a principal and an intermediary has been established and communicated to Government, the Gateway will check that a message is suitably authorised before passing it on. This will allow, for example, accountants to submit tax returns on behalf of their clients. A further level of registration allows the use of delegates, so that if the person who normally deals with your tax affairs is on leave, someone else can work on their behalf.

5. Supporting Implementers (UK GovTalk)

The aim of UK GovTalk (<http://www.govtalk.gov.uk/>) is to provide a library of XML schemas and support for those developing and using them. GovTalk uses and expands upon the set of standards defined in the e-GIF. The standards fall into two categories - the external standards such as XML and XML Schema, and internal standards and guidelines.

These cover areas such as the overall information architecture, a common set of schemas and schema components and guidelines for those developing schemas. These standards and schemas have two aims. The first is to provide a mechanism for passing data over the Internet. The second is to provide guidance for the specification of new systems. Currently, the vast majority of public sector systems have no support for XML. They therefore need some form of interface to support the common GovTalk standards. At some point in the future, these systems will be replaced, and at that time, the GovTalk standards can be used as part of the procurement process.

The information architecture aims to define common definitions for re-usable information items. Thus, for example, a citizen's name should always use a common base definition, although in the short and medium term there might be variations depending on the restrictions imposed by current systems. This architecture is then represented as a set of complex data types in several XML schemas.

As well as these schema components, GovTalk defines some complete schemas,

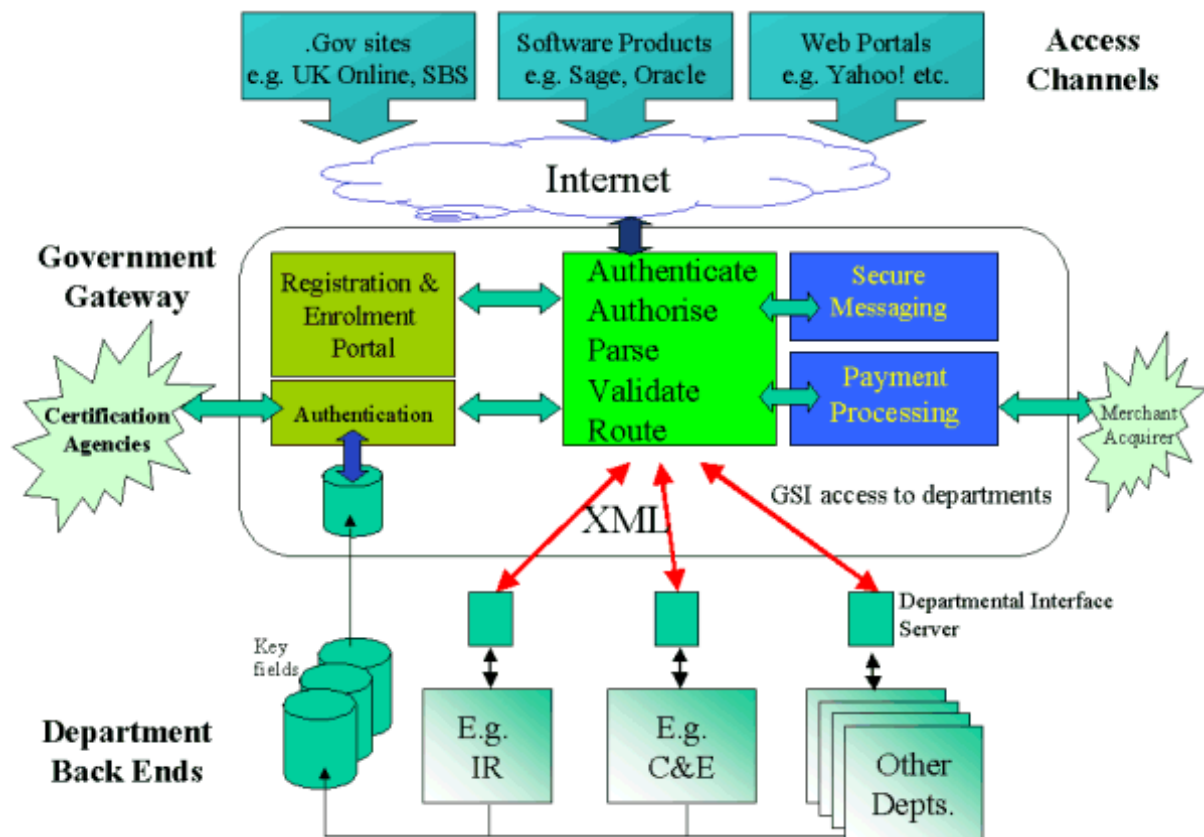
such as that used for the envelope within which all business messages are wrapped. This envelope is currently a custom solution, having been developed to meet the specific needs of GovTalk, but we are tracking the work of the W3C XML Protocol group and other bodies working in this area. The Office of the e-Envoy, which controls the e-GIF and GovTalk, is a member of the W3C.

As well as these standards, GovTalk supplies information and recommendations in areas such as how to develop schemas to be understandable to those less familiar with XML Schema. The schema guidelines have been developed in conjunction with the xml-dev mailing list members, who have been holding a long-running thread on this subject.

6. The Government Gateway

The Government Gateway is the key to making government look like a single entity to the business and citizen. It does this by allowing an individual to use a single credential (user name and password or digital certificate) to identify himself or herself across the entire public sector and by providing a single secure mailbox in which all communications from government can be stored.

The Gateway comprises three main sections, known as Registration and Enrolment (R&E), the Transaction Engine (TE) and the Departmental Interface Server (DIS). We will look at each in turn.



R&E performs two functions. Firstly, it provides the facility for people to indicate their intention to use government services by registering their credential (registration) and indicating the services they wish to use (enrolment). It then provides a service to the Transaction Engine to authenticate the sender of a message, indicate that the sender has enrolled for the service they are trying to use and is authorised to send the specific message. As an example, the first of these might indicate that I am representing ABC Accountants, the second that I have registered to use the PAYE service as an agent and the third that I am sending a PAYE return on behalf of alphaXML, which has authorised me to do so on its behalf.

The main function of the Transaction Engine is to route incoming messages to their destinations. In doing this, it might have to route to multiple destinations, and, as discussed earlier, might have to filter the message differently for different destinations. The TE first looks at the message to identify its type, for example, that it is a VAT return. It will then ask the R&E system to authenticate the user and authorise the transaction. If this is successful, it will often be able to route the message based on its type. In this case, the message will be passed to Customs and Excise. In some cases, other information in the message envelope will indicate the

public sector organisations that are to receive the message, so, for example, a change of address can be sent to the relevant local authorities as well as central Government departments.

The third part of the Gateway infrastructure is the DIS. This sits separately from the Gateway, being in the public sector organisations themselves. This device has two aims. The first is to hide the complexity of the protocols required for reliable message delivery from legacy systems. The second is to perform whatever translations are necessary on the incoming XML stream to interface to legacy systems, then to translate the outputs of these systems into the agreed XML format.

This has been only a simple overview, and has ignored some of the complications such as checking the digital signatures on messages. The aim has been to give an overview of the architecture.

7. Issues and Lessons

Obviously, in such a large project being developed in challenging timescales, mistakes have been made and lessons learnt. These lessons are now being applied to later phases of the project.

We realised from the start that providing common interfaces to public sector organisations would be difficult for several reasons. Apart from the existence of some very old legacy systems with different data requirements, public sector organisations have their own automation targets to meet, and fitting into the early phases of a common infrastructure is not always the quickest, easiest or lowest risk way of achieving this. Then each organisation has different definitions of common data items such as 'Name' and 'Address'. As a simple example, the restrictions on names in the UK are very few, allowing, for example, the use of a URL as a single name (yes - it has been done!). However, an individual organisation will often insist on a forename and surname because that is their existing business rule and their systems require it. Inevitably, creating general purpose XML schemas means that very few data items can be made mandatory. The schema development guidelines therefore allow extensions by restriction for individual data elements.

I say that we realised this would be difficult, but even so, perhaps we underestimated

the variations and the time it would take to get agreement. This has led to the first services available using the full infrastructure being those targeted at individual departments. The more 'joined-up' aspects of the project will follow later.

We have also found that we need to take more account of international standardisation efforts that might conflict with our information architecture. For example, we might define a model for an address, but what if a standard such as the Extensible Business Reporting Language (XBRL), which uses a completely different model for describing data through schemas, should become widely used? Clearly we should not exclude its use for purist architectural reasons. So these circumstances need to be taken into account.

There are also issues surrounding the use of draft standards. Currently the only standardised way of describing XML data is through the use of DTDs. However, we believed that these not only did not achieve all that we wanted in terms of data definition and reuse, but would also lead us down a blind alley. We therefore standardised on the April 2000 draft of XML Schema. Tool support for this draft is not strong, and that for the October 2000 candidate release has overtaken it. Should we move to this or wait for the full recommendation? How does this affect the infrastructure? How do we plan migration? These are all questions that have to be answered when dealing with a fast-moving technology.

These are perhaps the main issues. There are many others, some realised up front, others uncovered as we have progressed. The good news is that these issues are being considered and the architectures are proving robust. This is because, rather than design a grand fixed architecture, we have built in flexibility from the start. We would recommend the same approach to any large-scale integration project, whether in the public or private sector.

Acknowledgements

I am very grateful to the many people who have made this project so enjoyable and successful, and this paper possible. These include those at the office of the e-Envoy who commissioned and managed the project, the contributors to the work on XML schemas, the Microsoft team that developed the gateway, the people in the Government departments that became the lead users and the independent software

vendors that had the faith to develop products to submit documents through the Gateway.

Biography

Paul Spencer

alphaXML Ltd

United Kingdom

Paul Spencer - Paul has been specialising in XML since 1997, when he founded Boynings Consulting. He provides advice to the UK Government in its adoption of XML and has written the 'best practice guidelines' used for developing Government XML schemas. He also advises and develops schemas for several Government Departments. He is a member of the UK Government Interoperability and GovTalk working groups. He has worked with major commercial companies such as NatWest and Microsoft on their XML projects. Paul is the author of XML Design and Implementation and co-author of Beginning XML and Professional XSL (all Wrox Press). In 2000 he was a founder of alphaXML, of which he is CTO.