

Towards B2B Integration in Telecommunications with Semantic Web Services

Alistair Duke¹, Marc Richardson¹, Sam Watkins¹, Martin Roberts²,

¹ BT, Next Generation Web Research, Adastral Park, Martlesham Heath, Ipswich IP5 3RE, United Kingdom. {alistair.duke, marc.richardson, sam.watkins}@bt.com

² BT Exact, OSS Solution Design, Adastral Park, Martlesham Heath, Ipswich IP5 3RE, United Kingdom. martin.me.roberts@bt.com

Abstract. This paper describes BT Wholesale's B2B Gateway as an approach to provide Business-to-Business integration within the Telecommunications sector. Although the Gateway provides increased efficiency over separate systems, the process to allow business partners to integrate is lengthy and costly. The application of Semantic Web Services will ease the integration process. The Web Services Modelling Ontology is described and applied to the Gateway. The paper presents the initial requirements of the Gateway upon WSMO and proposes how WSMO could provide further benefit in the future.

1 Introduction

Today's telecommunications industry is facing many challenges. After many years of high growth and profit the last few years has seen rapidly falling prices and increasingly intense competition. Operators (and in particular, the large incumbents) have realised that they must radically transform the way they do business in order to reduce costs and remain competitive.

Currently the industry suffers from high manpower costs due to a lack of automation, poor time-to-market due to inflexible business processes and customer service which has suffered due to a lack of integrated support systems. On the other hand, customers are demanding integrated services, tailored to their specific needs. The market is becoming increasingly federated due both to regulatory pressures and to companies' attempts to catch market opportunities with tailored, bundled services. In this market, the number of Business-to-Business (B2B) relationships between telecommunications companies and specialist content and service providers has dramatically increased.

The current recession, changing market and new technology has led many companies to radically rethink the way they operate. They have realised that the new environment requires tighter management of processes and the eradication of bureaucracy and duplication of effort and systems. These requirements can be met by B2B integration where companies expose interfaces to their processes and systems, thus allowing their partners to integrate. This process, however, can be time-consuming and costly and can even result in proposed services being commercially

unviable. These problems are partly due to the legal and commercial aspects of forming a partnership and partly due to a lack of automation from the technical perspective.

The EU DIP project¹ and its technical basis, the Web Services Modelling Ontology aims to address the integration problem through the adoption of Semantic Web Services. In this paper we present initial work in applying WSMO to a telecommunications integration platform as part of a DIP case study.

2 BT Wholesale's B2B Gateway

Traditionally, vertically integrated telecommunications companies such as BT have provided end-to-end services to customers using their own retail operations and their own hardware. Over recent years, these companies have worked hard to improve customer service and reduce costs through greater process efficiency and effectiveness. These efforts have been enhanced with the introduction of integrated Operational Support Systems (OSS). These can provide customers with end-to-end visibility of service delivery and assurance. The challenge in the new environment is to maintain these levels of efficiency and customer service even though there are multiple parties and organisations acting to deliver the service who inevitably have their own systems that cannot be directly integrated with those of others [1]. BT Wholesale's B2B Gateway is provided to Service Providers² to allow them to integrate their Operational Support Systems with those of BT. Without such a system the service provider would either need to manually coordinate with BT via a BT contact centre or operate a system separate to its own OSS that communicated with BT's – thus requiring information to be duplicated.

The B2B Gateway exposes an interface whose behaviour is a combination of transport technologies such as SOAP, security protocols such as SSL, messaging middleware such as ebXML and the process behaviour of back end systems.

Over the last 10 or so years BT Wholesale has been involved in a number of B2B solutions. Earliest examples were based on Electronic Data Interchange (EDI) solutions for delivering billing information. The take up of these solutions was very low and it was not until the need to offer interfaces to other licensed operators for regulated products that their use increased.

About 5 years ago the use of XML based content was demonstrated. The content for this demonstration was based on a XML version of an American EDI library. This was publicly available and utilised an early version of a schema language known as SOX (Schema for Object-oriented XML). The strength of the library and the schema language it used was the ability to build simple extensions to support industry specific solutions.

Over that last four years a number of real solutions have been deployed by BT Wholesale using XML based on this early work. The most successful has been the DSL broadband interface that has been used to capture more than 1 million orders

¹ <http://dip.semanticweb.org>

² A service provider in this context is the organisation who has the relationship with the end customer.

With the publication of the XML Schema³ (XSD) standard there have been attempts to move the current XML interfaces away from SOX to XSD. The move has not been achieved for any live solution. For new solutions an XSD translation of the SOX library has been used or new XML schemas created where no library solution was available. The advantage of moving to the XSD format has been the availability of tools and the increased possibility of integrating with newer transport standards such as Web Services.

Currently the process involved in granting access to the Gateway for a new service provider is lengthy and complex. It commences with a commutation phase where partners assess their technical suitability, receive documentation and consider the level of fit with their existing OSS. A development phase follows this during which support is provided by BT. During the testing phase, the partner is given access to a test environment provided by BT where they can test the validity of their messages and their transport and security mechanisms. Firewalls, proxies, etc. must be configured by both parties to ensure that communication can occur. Once the testing phase is complete and documented the partner can move to a pilot phase where terms must first be agreed regarding volumes, frequency and support arrangements before access is given to the live system. Transactions are monitored during the pilot phase to ensure validity.

The process can take several months from start to finish. Any approach that can reduce development time, improve the quality of development through enhanced understanding and as a result avoid significant problems during the testing and pilot phases will naturally save BT and its partners significant time and money.

The Gateway currently exposes a number of interfaces concerned with service fulfilment and assurance. These are generally concerned with regulated services such as broadband access. The interfaces allow Service Providers to order and cease broadband lines on behalf of their customers, manage faults (i.e. raise and manage faults, request, confirm and cancel repair appointments and receive status fault status notifications) and carry out diagnostics (i.e. request tests and handle the response to these). In this paper, the application of Semantic Web Services to the Broadband Diagnostics interface is examined.

2.1 Broadband Diagnostics

As part of its OSS process, a Service Provider may wish to raise a test on the BT network. This is typically due to a problem that has been reported by one of its customers. The Service Providers OSS should collect the necessary information from the customer and assuming that the problem cannot be resolved internally issue a request via the B2B Gateway.

Interactions are implemented through the exchange of business documents, sent as messages. These interactions are known as transactions. The Gateway currently uses ebXML Business Process Specification Schema⁴ to model the sequencing of these transactions into a collaboration. The Broadband Diagnostics interface has only two

³ <http://www.w3.org/XML/Schema>

⁴ <http://www.ebxml.eu.org/process.htm>

transactions. These are 'RequestTest' and 'NotifyOfTestCompleted'. 'RequestTest' is a 'RequestResponse' transaction which means that a response to the test request is expected. This indicates whether the test has been accepted or rejected. It may be rejected if, for example, the Service Provider is requesting a test on a circuit that does not belong to them. The 'NotifyOfTestCompleted' is a 'Notification' transaction. This is a single message that is sent following the completion of an accepted test. It describes the results of the test. The sequence diagram for the collaboration in the case of an accepted test is shown in Figure 1.

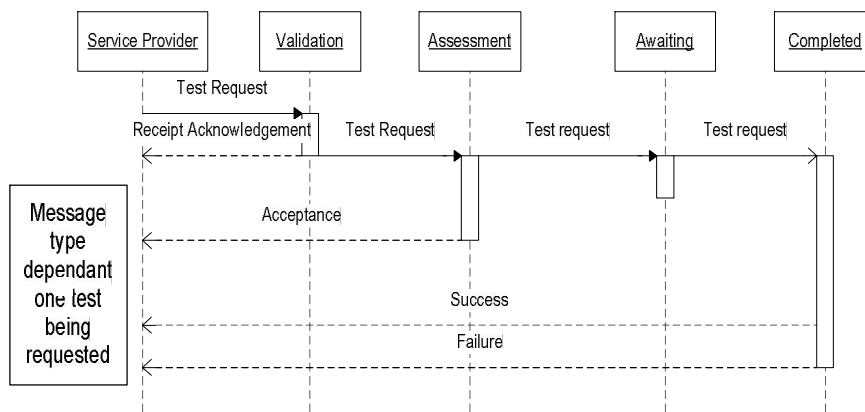


Figure 1. Sequence Diagram for Broadband Diagnostics Collaboration.

2.2 The B2B Gateway Architecture

The B2B Gateway, in common with most B2B interfaces has three separate elements. The two internal systems of the respective organisations that need to communicate and the interface that they will use to do this. This usually involves both systems translating their internal application view of data and process into the interface view of the problem. Depending upon who produces the interface definition the amount of translation involved can be either very small or almost impossible to achieve without development effort.

The Gateway architecture can be represented as shown in Figure 2. The Service Provider's OSS is able to generate a call to request a test. In order to pass this on to the B2B Gateway, it must first be adapted to enable it to be understood. The adaptation process has two key elements. Firstly the test call must be represented as a business message that will be understood by the gateway as a valid message given the current state of the transaction i.e. it must be represented as a TestRequest message

which is the initial interaction of the 'RequestTest' transaction. Secondly, the business message must be wrapped within the protocol envelope i.e. ebXML Messaging.

A message received by the B2B Gateway must also be adapted before it can be processed by the BT Wholesale OSS. This adaptation is effectively the reverse of the previous one. A message handler checks an incoming message for validity according to the message protocol and assuming it is valid unwraps the business message from the protocol. The business message is then checked for validity according to the current state of the transaction. Assuming this it valid, the message can be adapted into a call to the BT Wholesale OSS Test and Diagnostics system.

The system is asynchronous so 'Time To Perform' values are associated with the transactions. After processing the test request, the Test and Diagnostics system responds to the Service Provider via the B2B Gateway. The process is the exact reverse of the original request.

Generating the adapter between OSS calls and valid B2B Gateway messages is one of the key challenges of the integration process. The Web Services Modelling Ontology aims to significantly simplify this integration process. The next section briefly introduces WSMO and describes how the Broadband Diagnostics interface could be modelled using it.




Figure 2. B2B Gateway Architecture.

3 The Web Services Modelling Ontology

The Web Service Modeling Ontology⁵ along with its related language (WSML) and execution environment (WSMX) presents a complete framework for Semantic Web Services, combining Semantic Web and Web Service technologies. WSMO is an ontology and a conceptual model for the description Semantic Web Services. It is

⁵ <http://www.wsmo.org>

derived from and based on the Web Service Modeling Framework WSMF [2]. WSMO has a number of features (either defined or under development) that makes it appropriate (versus related approaches such as OWL-S) for application to the case study described in this paper. A detailed comparison of WSMO and OWL-S is given in [3]. WSMO features such as scalable mediation between loosely coupled elements and support for multiple groundings for processes are not well supported in OWL-S. Despite WSMO's advantages, it is the opinion of the authors that a complete departure from the W3C recommendation: OWL should be avoided. Although WSMO claims to be able to handle other ontology languages by providing a meta-level ontology, tool support for import and export has yet to be provided.

WSMO is based on two fundamental design principles:

- Strong *de-coupling* of the various components that realize an application
- Strong *mediation* to enable anybody to speak to everybody

The Web Service Modeling Framework consists of four elements. These are Ontologies which describe the domain in a machine processable way and provide the formal semantics to the information used by other WSMO elements; Goals, which are used to define a certain aim which will be achieved by using a Web Service, or the specific objectives that a client may have when consulting a Web Service; Web Services, which provide the functionality to process data or changing states in the physical world and Mediators, which are used to map between WSMO elements if they are based on different technologies, e.g. input and output data structures, business logics, message exchange protocols, etc.

The explicit relationship between services and ontologies is the key element for Semantic Web Services and WSMO. It is envisaged that this will enable:

- **Improved service discovery**

Semantic Web search technology allows users to search on ontological concepts rather than by keyword. This would allow users (and indeed computers) to find the most appropriate services more quickly and then narrow down their search via more expressive queries if required.

- **Re-use of service interfaces in different products / settings**

Services that are described semantically can more easily be discovered, understood and applied thus reducing the need to create new services that serve the same purpose. This could also be used in a strategy to reduce complexity i.e. remove services / interfaces that exactly repeat the function of other services but are described slightly differently.

- **Simpler change management**

Changes to models and services are inevitable over time. The key thing is to reduce the knock-on effect of change or at least manage it. A Semantic approach will significantly reduce the overhead and simplify the process. For example, when a proposed change is made to a data element, those services or interfaces that employ that data in some way can be dynamically discovered and appropriate action taken e.g. to contact the owner of the service with details of the proposed change.

- **A browseable, searchable knowledge base for developers (and others)**

In tandem with the example given above for simpler change management, semantically described services and ontologies would enable a knowledge base to be constructed. This would allow developers and solution providers to perform queries

relating to the data and processes they were concerned with, for example to determine the origin piece of data or its destination.

- **Semi-automatic service composition**

Given a high level goal which we wish a service or set of services to achieve, expressed in terms of an ontology, it should be possible to carry out decomposition into components parts and then match these components with appropriate services. The level of automation possible is a matter for ongoing research. Initial practical results are likely to provide users with a set of candidate services that might satisfy their needs. They are then left to decide between these services and oversee the composition required in order to satisfy the goal.

- **Mediation between the data and process requirements of component services**

Often there is need for two or more services to interact even though their communication requirements are semantically the same but syntactically different (they may require different message exchange patterns or different data formats). In this case it should be possible to automatically construct a translation between message data elements that allows the services to communicate. This process is known as mediation. It relies upon the mappings of messages and data elements to an ontology allowing semantic equivalence to be inferred.

- **Enterprise Information Integration**

As the name suggests, the Semantic Web is based upon web technology. This can afford universal (or at least enterprise wide) access to semantic descriptions of services (or information). One advantage of this is the ability to answer complex queries without having to consider how to access the various systems where the data required for the answer is held. For example, suppose there is a requirement to determine the number of customers within a particular postcode who spend more than £100 per quarter. If that information is held within one database and the person asking has access to it and knows how to query it then an answer could readily be obtained. Of course the situation is more complex if multiple databases hold the answer and access and a query interface have to be determined. The humans involved have some work to do in locating the data and processing it in the required way. A semantic approach, however, allows a single query to be made via a unifying ontology. [4]

3.1 Describing The Diagnostics Interface Using WSMO

This section describes a WSMO service description for the *testRequest* message.

WSMO service descriptions are separated into two parts – a functional part, describing what the service can actually achieve (service capability) and a non-functional part, providing additional information about the service (non-functional and quality of service properties).

Non-functional Properties

The non-functional properties for *testRequest* use the Dublin Core elements. These can be extended to include, for example, quality of service information. The property of most interest here is The property of most interest here is `dc:subject`, which

allows the service to be attributed to a class. Here we use an existing telecoms industry-wide process framework: eTOM (enhanced Telecoms Operations Map)[5] to provide a domain ontology. The value of the property is: eTOM:EvaluateAndQualifyProblem. This allows a direct link to be made between the service and a specific class in the eTOM ontology. The eTOM prefix is an XML namespace which is shorthand for the location where the eTOM ontology is defined. 'Evaluate & Qualify Problem' is an eTOM category in the 'Service Problem Management' process grouping. If appropriate, several eTOM categories can be attributed to a service using the subject property.

```

nonFunctionalProperties
  dc:title hasValue "Test Request"
  dc:subject hasValue "
    eTOM:EvaluateAndQualifyProblem"
  dc:description hasValue "Initiate a request for a
    Test"
  dc:contributor hasValue {
    <"http://www.btwholesale.com "> }
  dc:date hasValue "2004-12-10"
  dc:type hasValue
    <"http://www.wsmo.org/2004/d2/v1.0/#services">
  dc:format hasValue "text/plain"
  dc:language hasValue "en-US"
endNonFunctionalProperties

```

Functional Properties

The *testRequest* service could be described as having the capability that it initiates requests for diagnostic tests to be carried out over BT's network. In WSMO, modelling preconditions and postconditions provide the conditions over the information space that must hold before and after the execution of the service. Capability assumptions and effects also define conditions but describe the state of the world instead. As part of the modelling of WSMO services it is therefore necessary to identify what are the expectations and the influences of this service on the information space and the state of the world.

The inputs to *testRequest* are detailed in Table 1.

Input	Description
partyID	The identifier for the Service Provider who is placing the Request for a diagnostic test
refNumber	The reference that identifies a particular test for the Service Provider
testRequestDateTime	The date and time the test request was sent.
testCategory	The type of test requested

Table 1. Inputs to TestResquest transaction

Preconditions of this service would be is that the *partyID*, *refNumber* and *testCategory* are all valid. In WSMO this is modelled with Axioms:

```
?refNum memberOf refNumber and validRef(?refNum) .
```

```
?testCat memberOf testCategory and  
validTestCat(?testCat) .
```

```
?parID memberOf partyID and validpartyID(?parID)
```

where *refNumber*, *testCategory* and *partyID* are predefined concepts and *validRef()*, *validTestCat()* and *validPartyID* are predefined axioms in the same ontology. *refNum*, *testCat* and *parID* are variables representing the input to the *testRequest* service.

The output of the service will be the results of the test, so the postcondition would be a valid test result:

```
?testRes memberOf testResult and validTestRes(testRes)
```

where *testResult* is a pre-defined concept and *validTestRes()* is a pre-defined axiom in the same ontology. *testRes* is a variable representing the output of the *testRequest* service.

To enable definition of assumptions and effects of the *testRequest* service the state of the world should be modelled. In this scenario a *testRequest* could have one of four states:

Ready – The test request is ready for submission

Awaiting – The test request has been accepted but the test has not been completed yet.

Completed – The test has been completed

Failed – The test request was not accepted and no test was carried out.

The states are modelled through variables, and the values of these indicate the state of the world at a given time. The state of the world in which *testRequest* can operate (assumption) is one where the variable *TRstate* has value “*Ready*”. The execution of *testRequest* changes the world state (effect) so that *TRState* value changes to “*Completed*”. The WSMO axiom defining this assumption can be defined with the following expression:

```
? WState memberOf WorldState and
```

```
? WState [TRState hasValue “Ready”]
```

where *WorldState* is a predefined concept having *TRState* as an attribute. The effect is similarly defined:

```
? WState memberOf WorldState and
```

? WState [TRState **hasValue** "Completed"].

Thus WSMO can be used to describe a rich set of properties that the service has including its capability, its data requirements for input and output and its effect of the information space and state of the world.

3.2 A WSMO-based B2B Architecture

Describing the Broadband Diagnostics interface using WSMO allows the B2B architecture picture to be redrawn as shown in Figure 3.

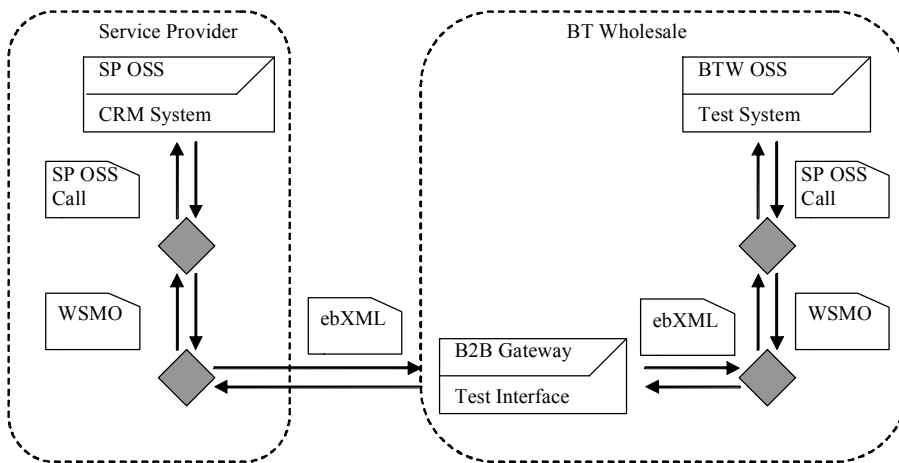


Figure. 3. A WSMO based B2B Architecture.

Initially it may appear that the WSMO architecture is more complex since one extra adapter has been added to each party's side of the architecture. OSS calls are mapped to WSMO before being mapped to ebXML messages. However, from the Service Provider's perspective, the WSMO representation of the Gateway already exists – BT Wholesale has provided it. The Service Provider is required to map its calls to WSMO and then transform this representation into the WSMO representation of the Diagnostics interface. This still sounds like two major adaptation steps, however, the point is that the first of these could be automated and the second is at the ontological level and should be easier much easier to perform than the equivalent 'syntactic' transformation that is required with the current B2B gateway. The reasons behind this and approaches for automation will be identified in section 4.

From BT Wholesale's perspective, initially, more work is required than with the current Gateway. Firstly, it must represent its interface using WSMO. This requires developing ontologies (either from scratch or by building upon any industry ontology work such as that in the TeleManagement Forum [5]). The ontologies must represent the complete data requirements of the interface as well as the process of collaboration. However, following this effort it is to be expected that ongoing effort will be reduced

since Service Providers will require less support and should be able to integrate more quickly with fewer errors.

Whether the WSMO adaptation within BT Wholesale's part of the architecture is required is an interesting point. Initially, this seems unnecessary since BT should fully understand how to integrate here and should only need to do it once. However, in practice both of these statements are rarely true: often the people responsible for internal systems are different from those for external integration. Also, there is no reason why the internal systems should be static. Updates are always required to existing systems and software vendors may change resulting in complete platform changes. In this environment a semantically described internal interface makes more sense.

4 Mapping Messages

Designing interface content appears at first glance to be very simple. Simply gather up the dataset that you know your internal systems require and package them into suitable format using straightforward and unambiguous field names. XML brings only one significant advantage over the original fixed length records – any examples produced during design will look exactly like the final solution. This means that early examples can be 'war gamed' to verify a solution before any code is written. However, beyond this XML still requires about the same amount of explanation and documentation as the earlier cryptic file formats. The reason for this is purely one of semantics. XML is only an 'improved syntax'. When humans read XML they bring their own semantic understanding to bear and this process initially makes it appear that XML is an improvement as it is supposed to be 'self-describing'.

This apparent self-describing quality of XML has given rise to a huge number of misunderstandings when interfaces are implemented. For example is the contact for a Service Provider the person who is managing the transaction or the customer that the Service Provider is supporting? There have been examples where the same set of fields have been used to convey both interpretations. This means that BT Wholesale have difficulty knowing if the contact information provided is suitable for them to use. Another inherent problem with XML is the apparent ease with which anyone can produce it. This inevitably means consistency is threatened.

4.1 XML 'Semantics'

There are a number of issues with representing and determining the semantics of XML including:

Structure: The hierarchical design of XML means that the placement within the structure can hold the key to the semantics of a particular data item. For example the following are all semantically equivalent but syntactically different:

1. `<ServiceProviderReference>123wqrh10</ServiceProviderReference>`

2. `<ServiceProvider>
 <Reference>123wqrh10</Reference>
 </ServiceProvider>`
3. `<ServiceProvider Reference="123wqrh10"/>`

Meaning: The meaning of a field may change dependant upon who is looking at the message. Examples of this are the typical ‘Our Reference’- ‘Your Reference’ fields. The problem of using these ‘relative’ terms is that the semantics can only be interpreted if the context, i.e. who sent the message, is understood at the time of interpretation.

Inter-field links: The traditional schema languages that are used to describe XML have no mechanism to describe inter-field links. These are where a value in one field indicates that a particular set of fields should exist or adopt certain values in a different part of the message.

This means that it is possible to create ‘valid’ messages that parse successfully through a schema validator, yet they will fail when applied to the receiving system. This can force the XML designer to adopt extra layers of hierarchy to show these links and this can mean obscuring the essence of the messages meaning.

In order to enable successful understanding of what the B2B Gateway interfaces actually mean a great deal of documentation has been required. It is ironic that in order to get computer systems to communicate successfully that you first need to communicate successfully with fellow humans. This process usually gives rise to ambiguities as to the meaning of fields. The actual process of producing documentation and associated tools such as example generators and validators actually means that the designer has to explain the meaning of each field. To help this a technique of adding metadata to the underlying schemas has been used. This metadata was defined by the ITU in a standard called tML (Telecommunications Markup Language)⁶. Adding the metadata does help document the content but it slows the process of designing interfaces down and can sometimes introduce incorrect definitions as the XML designer is not usually the person who really understands the underlying data and its associated meaning.

4.2 Achieving successful integration

To achieve a successful integration of two systems, they must both understand the process and data related to the problem space. Increasing the understanding of particular messages and data items through the use of XML syntax helps, but it does not remove the need for humans to do the mapping from the internal systems to the external interfaces.

This mapping involves the developers/designers in a semantic exercise that is often underestimated in its complexity and the time it can take to code a solution. This time and complexity is reduced significantly if the domain of the problem is understood by both parties and the systems using the interface have been built with the interface in

⁶ <http://www.atis.org/>

mind or with at least an understanding of the domain involved. For example, if the Order system being used to produce orders has the ability to capture the required data in a manner that is similar to the requirements of the interface the chances of a successful integration are increased.

4.3 How could having a Semantic layer help?

If you take two systems and imagine that they both have a clear semantic definition of their capabilities, namely the processes and data they support, being able to access the ability of the systems to integrate should be fairly straight forward. To date the issue has been that this Semantic layer has been in the form of document (usually Word and Excel) and is aimed purely at a human readership. If these semantics could be expressed in a computer readable form, it would be possible to do a gap analysis of the data, but also it should be possible to automatically build the mediation layer between the systems. This could mean that the interface between the systems is actually bespoke to them rather than being published as a generic interface.

This semantic layer would have to encompass not only a one dimensional understanding of the data at face value – i.e. this is a date, but also be able to understand that this date is the ‘Customer Required by Date’. Above this it would have to understand that this is the date when the product associated with the Order is to be provided. To develop this level of semantics the Semantic Layer needs to be able to express understanding not just about the messages themselves but also the underlying concepts associated with them, such as products, services, bills, invoices, etc.

With this deep understanding it becomes increasingly possible to change the way in which B2B solution actually work. From transactional model it is possible to move to a model where the exchange of relevant information and state when appropriate is the norm. The computer systems would be able to derive which bits of information they require from each other. This could mean that the systems ordering a product could be asked for extra information which is not known at the start. The fact that this new information is ‘understood’ by the Semantic layer would allow the system being asked for this information to assess if it has it or if it needs to ask a human to provide it.

Although the above utopian dream is goal worth expressing it is also worth looking at what can be achieved today with the Semantic tools available. Assuming a given process with a given set of data it should be possible to express a set of semantics that describe this finite world. By expressing an ontology of these semantics it is possible to map a given systems data into the ontology. This means that if the interface needs were expressed in terms of the initial ontology then the mapping from the system to the interface could be automatically generated.

4.4 Mapping to an Ontology

It is possible to map between XML and an ontological representation such as RDF using XSLT. However, this is problematic since two unrelated transformations – one

from XML to the ontological layer and another in the opposite direction are required. In addition, because of the graph structure of RDF, there is no canonical serialisation that can ensure sequence information is maintained. A promising recent approach [6] which overcomes these issues is perhaps the most appropriate in this scenario. A bi-directional mapping is enabled through the naming of schema components which preserves the scope of elements and attributes in XML. In additions, the DIP project has considered the state-of-the-art in this area [7] and will build a WSMO mediation component that can be applied to the B2B Gateway.

5 Looking Ahead

As described in section 2, the common scenario for B2B integration in Telecommunications involves long-standing partnerships and agreements. Technical integration takes place following a legal and commercial process where service levels, price and other agreements are reached. In this scenario the need for dynamic, run-time discovery and composition is not obvious. Even at design-time, discovery is not currently required since a relatively small number of interfaces are available and these are well-known and documented – this explains the large focus on mediation requirements in this paper which most definitely is required at design-time.

This scenario is expected to change dramatically in the near future. The long-standing partnerships are unlikely to disappear overnight but shorter-term, ad-hoc collaborations are expected to emerge. For both of these forms of relationship there will be a greater reliance on dynamic integration. Organisations will expose more-and-more of their business interfaces to both their long-term partners and their customers but also to anyone else who wants to do business (of course with the appropriate restrictions). This trend is due to both regulatory pressure (e.g. as seen in the UK with ‘local-loop unbundling’) and the need to make supply chains more efficient allowing cost reductions and greater agility in service delivery. In this more open environment the level of integration, of course, increases; as do set-up and management costs. There are severe doubts as to whether a static integration approach is sustainable under these circumstances. Even if design-time integration is still heavily relied upon, the ability to discover alternative functionality, perhaps for resilience or where factors such as price and availability are changing quickly, will be required.

How does the B2B Gateway fit into this changed environment? Many more service, wholesale and content providers are expected to expose their interfaces for integration. It is here that technologies such as WSMO have real value since the initial effort required in creating ontologies, describing interfaces semantically and relating the two together is now much less than the total integration effort. The ontology created by BT Wholesale to describe their interfaces will perhaps be adopted by other smaller suppliers since they realise that service providers who are already using it can more easily integrate with them if they do so. Other ontologies are created by other dominant players in other markets such as global content providers. Here mediation between the ontologies is required for integration. Discovery in this environment leads to a dramatic change in business relationships. Today, companies ask ‘Who are

my partners?' then 'How do I integrate with them?'. Tomorrow, companies may ask 'What services can I discover?' then 'How do I partner with their owners?'.

6 Conclusion

This paper has presented an existing approach for integration within the Telecommunication sector – the BT Wholesale B2B Gateway. Although providing increased efficiency when compared to separate OSS systems the process to enable BTW's partners to use the Gateway is long and costly. It is proposed that Semantic Web Services could ease this integration process. The Web Services Modelling Ontology is presented together with a description of how it could be applied to the Gateway today and how it might provide further benefit in the future. Initial benefits are expected in the mediation space with discovery and composition aspects undergoing consideration following this. The paper describes initial work on a case study within the DIP project. The application of WSMO and associated tools that emerge to this case study will continue to be explored within the project.

References

1. Evans, D., Milham, D., O'Sullivan, E., Roberts, M.: Electronic Gateways—Forging the Links in Communications Services Value Chains. *The Journal of The Communications Network* Volume 1 Part 1 (2002)
2. Fensel, D., Bussler, C.: The Web Service Modeling Framework WSMF. *Electronic Commerce Research and Applications*, Vol. 1, Issue 2, Elsevier Science B.V. (2002)
3. Lara, R., Roman, D., Polleres, A., Fensel, D. A Conceptual Comparison of WSMO and OWL-S. In *Proceedings of the European Conference on Web Services (ECOWS 2004)*, 2004.
4. Duke, A. Davies, J., Richardson, M. Kings, N.: A Semantic Service Orientated Architecture for the Telecommunications Industry. *INTELLCOMM (2004)* 236-245
5. TeleManagement Forum, Enhanced Telecom Operations Map (eTOM) data sheet. Available on the web at: <http://www.tmforum.org/>
6. Battle, S.: Round-tripping between XML and RDF. *ISWC04 Poster Presentation (2004)*
7. Grimm, S. et al.: Service Mediation: State-of-the-art and Requirements Analysis, DIP Deliverable 5.1 (2004)