ORACLE
INSURANCE

Oracle Health Insurance Back Office

Back Office Service Layer

User Manual

Version 1.5

Part number: E54856_01

May 26, 2014

Internal code CDO 14826

CHANGE HISTORY

Release	Version	Changes
10.12.2.0.0	1.3	• Added technical principles at the consumer webservices
10.13.3.3.0	1.4	• Extended with starting points that apply to the find services with mostly optional arguments.
10.14.1.0.0	1.5	• Included paragraph on 'write' services from 'Custom Development for Oracle Health Insurance Back Office' • Added write services to 'common usage' paragraph.

RELATED DOCUMENTS

A reference in the text (doc[x]) is a reference to another document about a subject that is related to this document. Below is a list of related documents:

- **Doc[1]**: OHI Back Office Service Layer Installation & Configuration manual (internal code CTA 13651)

Contents

1 Introduction

With release 2011.02 of the OHI Back Office application a first version of the 'Service Layer' has been released.

The Service Layer is an optional component that in the long term should offer all mainstream services to retrieve and manipulate the core OHI Back Office 'fact' data (so the mass data which is somehow member related; typically this is the data that does not fall in the category of configuration or 'setup' data). The Service Layer in this meaning 'provides' a set of 'provider services', services that can be 'consumed' by other applications.

But in a later release also a set of calling out web services, consumer web services, has been rebuilt. These web services are also part of the Service Layer, although they are less directly used, more indirectly in application functionality.

The rest of this document focuses on these provider and consumer services.

The Service Layer is targeted to ease integration in a Service Oriented Architecture (SOA) environment. However, in order to leverage investments, knowledge and experience and to benefit throughput, the provider services are also offered as plsql services in the database.

This document describes the generic technical details regarding the service layer, how to use it and what it is intended for.

Because the provider services are offered both as SOAP and as PL/SQL implementation this documentation normally does not distinguish between both implementations. This is because the SOAP implementation is built on top of the PL/SQL implementation.

The functional specific details (additional to their functional behaviour documented in other documentation like the Online Help) per provider web service are described in this document in an appendix per web service. The consumer services are more used as integral part of application functionality and are currently not separately documented. They are currently implementations that implement (part of) the functionality as offered by provider web services from external parties.

Most of the chapters in this document focus on the provider web service functionality.

For information regarding installation and configuration of the OHI Back Office service layer components please use Doc[1]. That document also contains a brief description of the architectural setup of the service layer.

2 Generic principles provider services

This paragraph describes a number of generic principles that apply to a subset or all Service Layer provider services and which are good for understanding best how the services have been developed. Goal is to provide information on how to use these services best.

2.1 Design principles

A number of design principles are being used in the setup of the provider services part of the Service Layer. They are listed below and may have influence for how to use the Service Layer provider services.

For the OHI Back Office Service Layer provider web services the following principles are used (during the initial transition period where new and previous services exist next to each other there may be temporary exceptions):

1. Terminology, terms and documentation will be in the English language.

   Rationale: the service layer is typically used by developers which are required to know English because most tooling documentation is also only available in English.

2. Terminology, terms, structure and contents of the service definitions will be aligned as much as possible over the different Oracle Health Insurance product lines.

   Rationale: to ease implementing functionality which crosses boundaries of different product lines it helps when the same terms, etc. are used. However, strict (technical) dependencies are prohibited to prevent complicating maintenance dependencies, so differences can occur.

3. A provider web service layer as well as a similar plsql service layer will be offered.

   Rationale: the first layer is offered to access the application from other heterogeneous applications, the second layer for accessing the OHI Back Office application from custom code within the same database (in such situation it would be a complete overkill to use a web service interaction although it can offer a more loosely coupled application architecture).

4. Existing OHI Back Office web services based on older technologies, and developed for other reasons, will be replaced in the coming years with at least similar but usually more extensive and generic services functionality in the new services layer.

   Rationale: this will help in offering a uniform and consistent way of implementing services; it will ease management and reduce maintenance efforts that finally will benefit customer investments.

5. The service layer will focus on offering quite generic usable services instead of offering a very application (or localization) specific limited service.

   Rationale: by offering more generic services these can be used for all kind of different integration purposes instead of only for a very application specific interface.

2.2 Technical principles

The following technical principles are followed and may be of influence for how you realise your code to access the Service Layer provider web services.

1. SOAP 1.1 is used.

   Rationale: this is by far still the most widely used common standard and supported by almost all web service toolkit implementations.

2. Document style web services are used.

   Rationale: this makes the services widely usable because they are implementation and platform independent.

3. WebLogic Server will be used as the standard application server deployment platform.

   Rationale: this is a highly scalable, reliable and robust application server for deploying Java applications that offers a lot of out of the box functionality.

4. Security functionality will be externalized from the web service implementation unless Oracle standards require differently.

   Rationale: by externalizing authorization, authentication and encryption from the service implementation it is prevented that existing customer functionality conflicts with chosen implementation solutions. This means that customers can leverage their investments for as far as these can interact with a Java based WebLogic deployed application. Normally applying open standards Middleware to support these functionalities will be the way to go.

5. The service calls will be stateless.

   Rationale: to serve easy integration each call is stateless; there is no state to be remembered over more than one service call.

6. A single change call will contain one or more atomic transactions.

   Rationale: service calls that change data do this in one or more atomic transactions, which completely fail or succeed, to prevent inconsistent situations can arise.

7. Optionally a form of optimistic locking can be activated. Default no locking will occur.

   Rationale: because services are stateless and to prevent long outstanding locks blocking other transactions a form of optimistic locking can be activated to control concurrency impact. When at a certain moment data is read and the functionality requires that changes be made under the requirement that the earlier read data is not changed, this optimistic locking can be activated. When the service, which changes the data, detects that the data has been touched by a transaction that committed since the last read moment the change will be rolled back and an error will be returned.
   Of course always a record will be locked explicitly immediately before it is changed, to prevent hanging service calls due to outstanding changes on a record. When the actual lock fails an error message is returned and the transaction is rolled back.

8. SOAP error handling will be used to return functional and technical errors.

Rationale: this eases integration with development tooling which can leverage functionality based on SOAP faults and makes coding easier and less error prone.

9. Functional faults will support language dependent error messages.

   Rationale: although the services are in English the functional error messages returned will use the multi-language support as present in the OHI Back Office application; this to be able to return language specific messages based on the calling context.

10. Standard Java logging functionality will be offered for error, informative and debug level log messages.

    Rationale: by adhering to a common standard logging mechanism this will be easier to configure and use for system administrators who are experienced with Java based application management.

11. The user manual currently focuses on a 'standardized approach' for synchronous provider web services.

    Rationale: when this manual applies to other types of web services the manual will be adapted for this.

12. Additional standardized requirements will be implemented as much as possible through applying standardized technology.

    Rationale: goal is to focus on delivering functionality and be open for most architectural and infrastructural environments; this means that limited development time is not spent on developing proprietary solutions that can be implemented also through standard technology stack software. A number of requirements should be implemented through more or less standardized use of Oracle products. But customers may opt out for other choices. Versioning of services is an example of this.

13. Date and date time values are not expected to have a time zone component in them.

    Rationale: the current application does not support time zones and all date and (date)time values are expected to be expressed in the time zone as used within the database. It may be that time zone handling for (date)time values in the web services are added in a future version where these are always converted to a standard time zone.
    When time zones are passed in values it is expected a service bus or different proxy solution will remove this component in the date and (date)time values.

3 Generic usage aspects provider services

This chapter focuses on the generic aspects of the web services.

3.1 Common usage behaviour

For retrieving data several operations may be defined. These operations can be distinguished in 3 types for which certain behaviour rules apply.

- Find services
  The find routines can be used to find a set of occurrences given certain arguments. The arguments in itself do not always need to be existing values and may contain wildcards in a number of cases (see later for more generic principles for find services). If no data matches the criteria no response data is returned and no functional error message is given. Simply the fact that nothing is found should make clear that no data is found.

- Get services
  Routines, which implement 'get' functionality, expect identifiers as inputs that identify exactly one occurrence. If the identifiers do not identify an existing occurrence a functional fault will be returned that no data matches the criteria.

- Write services
  Used for selectively updating a single occurrence, for example a relation or provider contract.

3.2 Find Services – Usage notes

There are two main groups of find services. One set contains some required arguments that are selective and make sure a quite limited set of data is searched for. Some examples are find services for retrieving errors, coverage's or flex fields for a given claim line. Or finding claim lines for a specified member.

Another group of services are more generic find services with almost all arguments optional. These are meant to search for data based on quite different arguments. For these services it is important to take notice of the following starting points and usage requirements:

- When an argument is entered a value for, it is used as search condition. It is not possible to search for occurrences where that argument is explicitly empty (not entered, 'null' in database terminology).

- For string conditions on 'fact data', meaning arguments that do not reference a setup identifying value (like a country code or a name of a flex field), wildcards are allowed. Wildcards are specified using 'Oracle database' wildcards like '_' and '%'. A SQL 'like' will be implemented in such situations. Wildcard usage may result in expensive inefficient queries if used in indexed access paths by being but quite at the beginning of a string making it not being very selective.

- For string conditions usually the search is implemented case insensitive. Exceptions are selective conditions that can be implemented by an indexed access path being present without special functionality to enable efficient case insensitive searches. An example is postal code, this is case sensitive. Searching on a relation name is not case sensitive because special facilities enable efficient case insensitive searches.

- When an argument is used to identify a related selective condition, for example the 'external code type' for an 'external code', both the identifying argument as

well as the selective argument needs to be entered. Otherwise the service will fail with an error. Typically the uppercase value or the passed argument will be used as most of these arguments are by definition stored in uppercase internally (so you do not need to care about case sensitivity).

- When a condition is entered that identifies fact data, for example a country code or an external code type, no wildcard is allowed and existing value needs to be specified. The argument entered will itself be checked for existence before a search is started. The service call will fail if the value does not exist.

- There is no protection against misuse of the services by not specifying any selective arguments. It is expected from the implementing party that at least a selective argument is filled with a selective value. This means that for example searching for all relations with an address on number 24 or with a name containing somewhere 'ee'(using argument '%ee%') will result in expensive inefficient queries that on environments with quite some data may result in a time-out of the service call while the query keeps on executing on the database. It is the responsibility of the calling party to prevent such misuse.

- Given the previous bullet it is expected that realistic performance and usage tests are executed on a representative environment (with production like amounts of data and similar resources and load). This is imperative to prevent response time and load problems are only detected in production.

- As a bold measure to prevent very long running services queries it is possible to activate resource limits on the database for the database account that implements the service calls.

  A very simple example to prevent a call uses more than 50.000 block gets is shown below:

  ```
  create profile svl_prof_max_gets limit
  logical_reads_per_call 50000;
  alter user <svl_user> profile svl_prof_max_gets;
  alter system set resource_limit=true;
  ```

  The limit will be applied to newly logged on database sessions. ORA-02395 will be raised when the limit is exceeded. It is a bold way to prevent a connection pool will become exhausted.

3.3 Write Services

The current, second-generation of write services supports idempotent behavior, selective updates and processes time-valid data. The first service of this new breed is the WriteRelation service.

The existing first-generation 'write' services will be migrated to the new paradigm over time.

3.3.1 Idempotent behavior

'Write' services should have idempotent behavior to ensure that each subsequent call to a business service with the same data will return the same response and have the same effect as the first call.

Note that this requirement does not apply for 'Find' and 'Read' services, because the contents of the data base may have changed between subsequent service calls.

3.3.2 Selective updates

'Write' services are used both for inserting and updating data into the OHI Back Office data base. When updating, ´write´ services support selective updates to allow the caller to send a partial message. The advantage is that the calling application only needs to know a subset of the data which can be processed by the business message. For example a self-service application only needs to have very little data to allow an end user to update his residential address.

In the example below, the name and phone number are set for relation 1864856800:

```
<v11:Person>
  <v11:relationNumber>1864856800</v11:relationNumber>
  <v11:name>Bakker</v11:name>
  <v11:phoneNumber>06-51227410</v11:phoneNumber>
</v11:Person>
```

If we want to change the name, we can simply pass the new name:

```
<v11:Person>
  <v11:relationNumber>1864856800</v11:relationNumber>
  <v11:name>Slager</v11:name>
</v11:Person>
```

We can also wipe the contents of a column, for example the phone number:

```
 <v11:Person>
  <v11:relationNumber>1864856800</v11:relationNumber>
  <v11:phoneNumber></v11:phoneNumber>
</v11:Person>
```

3.3.2.1 'Zero' update

Leaving out a tag means that its related data is left untouched. This is the cornerstone for selective updates.

3.3.2.2 Use xsi:nil to remove existing values

With empty tags you can wipe lists and simple string values. If you want to wipe values which are enumerations or other data types, you should use the nil attribute, like below:

```
<v11:startDate xsi:nil="true"/>
```

Note that the xsi namespace should refer to http://www.w3.org/2001/XMLSchema-instance.

3.3.2.3 Lists

Lists can have 0 or more elements which together are enclosed in a list-tag, like in the example below:

```
<v11:Person>
  <v11:relationNumber>1864856800</v11:relationNumber>
  <v11:bankAccountList>
    <v11:bankAccount>
      <v11:accountNumber>
       NL42RABO0111750768
      </v11:accountNumber>
       <v11:bankRelationNumber>
        1525725800
      </v11:bankRelationNumber>
       <v11:bankAccountType>IBANAccount</v11:bankAccountType>
       <v11:countryCode>NL</v11:countryCode>
       <v11:currencyCode>EUR</v11:currencyCode>
    </v11:bankAccount>
  </v11:bankAccountList>
</v11:Person>
```

To support selective updates, lists are optional. If you do not want to update a list of details, just omit the list and its surrounding list tag:

```
<v11:Person>
  <v11:relationNumber>1527308300</v11:relationNumber>
  <!--<v11:bankAccountList/> -->
</v11:Person>
```

Likewise , if you want to delete the list, use an empty list tag:

```
<v11:Person>
  <v11:relationNumber>1527308300</v11:relationNumber>
  <v11:bankAccountList/>
</v11:Person>
```

Note: if you add a list to the request you must include all elements. You cannot add a list with a single element just to update the single element. In that case all other elements will be deleted.

3.3.2.4 Time-valid lists

Time-valid lists are used to create and update data with a start and end date, such as addresses, marital statuses etc.

They share some characteristics with ordinary lists:

- Time valid lists are optional: the list related data in the OHI Back Office database are not updated if you omit the list tag altogether.
- All list-related data in the OHI Back Office are deleted if you specify an empty list tag.

The difference is that you can use time-valid lists to update the current situation without removing past data.

Consider the following example to register that Peter is no longer married since 1 January 2013:

```
<v11:maritalStatusList>
  <v11:maritalStatus>
    <v11:startDate>2013-01-01</v11:startDate>
    <v11:maritalStatus>
     dissolved marriage / dissolved registered partnership
    </v11:maritalStatus>
  </v11:maritalStatus>
</v11:maritalStatusList>
```

This information is processed as follows:

- The start date of 1 January 2013 is used as a reference date.
- Peter's previous marital status record (married) is ended by 31 December 2012
- Peter's marital status records starting after the reference date are deleted (if they exist)
- A new marital status record to indicate Peter's current status is created with a start date of 1 January 2013.

3.3.2.5 Segmented time-valid lists

The mechanism described above is too coarse for processing time-valid information like addresses, since you may have different home and postal addresses at any point

in time.
This is solved with segmented time-valid lists: this means that the list is processed once for every segment, for example 'address type'.

Consider the following example where John's home address is updated:

```
<v11:addressList>
  <v11:address>
    <v11:startDate>2010-06-04</v11:startDate>
    <v11:addressType>Home</v11:addressType>
    <v11:street>Haverstraat</v11:street>
    <v11:houseNumber>41</v11:houseNumber>
    <v11:postalCode>3511NB</v11:postalCode>
    <v11:countryCode>NL</v11:countryCode>
  </v11:address>
</v11:addressList>
```

This information is processed as follows:

- The start date of 4 June 2010 is used a reference date for John's home addresses
- John's previous home address is ended at 3 June 2010
- John's home addresses starting after the reference date are deleted.
- A new home address is registered starting 4 June 2010
- John's postal addresses remain untouched.

We can update John's home addresses and postal addresses in one go:

```
 <v11:addressList>
  <v11:address>
    <v11:startDate>2010-06-04</v11:startDate>
    <v11:addressType>Home</v11:addressType>
    <v11:street>Haverstraat</v11:street>
    <v11:houseNumber>41</v11:houseNumber>
    <v11:postalCode>3511NB</v11:postalCode>
    <v11:countryCode>NL</v11:countryCode>
  </v11:address>
  <v11:address>
    <v11:startDate>2010-07-01</v11:startDate>
    <v11:addressType>Postal</v11:addressType>
    <v11:street>Postbus</v11:street>
    <v11:houseNumber>306</v11:houseNumber>
    <v11:postalCode>3300AH</v11:postalCode>
    <v11:countryCode>NL</v11:countryCode>
  </v11:address>
</v11:addressList>
```

Note that an empty address list will delete all John's addresses:

```
<v11:addressList>
</v11:addressList>
```

Note: segmentation is not necessarily restricted to a single element (like address type in this case)

Note: consult the functional specification to find out whether a time-valid list is segmented and which elements are used for segmentation.

3.3.2.6 XSD types for Write services.

When examining an XSD for a web service with 'Write' operations you will find that the complex types used by Write Services are prefixed with 'PX'.

You will also find that these complex types largely consist of optional elements. This is needed to support selective updates.

There is a downside to this optionality: if you leave out many elements when entering new data, your inbound XML will still validate correctly against the XSD.

However when sending the request, the OHI Back Office business rules come into play and raise exceptions if your data is incomplete or incorrect.

It would be too complex to document which business rules you may encounter.

Our advice for validating a client application using a 'Write' service would be to always include various tests with new data.

3.4 Call standards

For each service a calling context and a returning context must be provided to call a service routine. The calling context specifies behaviour and the returning context provides feedback about the call.

These two 'contexts' are described by referencing the plsql definitions for these contexts (these are 'published' identically in the Java layer).

For all service calls there are 2 standardized SVL object types which need to be passed as 2 separate parameters to each service call.

- ✓ SVL_CALL_CONTEXT_TP - input parameter set to pass call data
- ✓ SVL_RETURN_CONTEXT_TP - output parameter set to pass return data

These are each described separately but they are related to each other when optimistic locking functionality should be implemented. In this latter situation part of the return context of a preceding call is input for the next call context.

3.4.1 Call Context

The (plsql) definition of the call context is as follows:

```
create or replace force type svl_call_context_tp
as object
(
  user_context                svl_user_context_tp
, enforce_consistent_read     svl_yes_no_tp
, enforce_unchanged_since_scn number
)
```

The call context is used to define the behaviour of the called service. A number of settings can be provided to the call.

3.4.1.1 User_context

The user_context setting is used to pass the OHI Back Office known and active username which should identify the user that executes the action.

3.4.1.2 Enforce_consistent_read

The setting enforce_consistent_read is used to ensure that in a service call all eligible data retrieved in and returned by that call is consistent, in the sense that it is all not changed (and actually committed) since the call started (so there may not be a committed change on the retrieved data by another session, since the retrieve operation started, to prevent sequential selects in the retrieve call retrieve an inconsistent situation; with other words, the data returned by the call is as it was at the moment of that call, it is a 'stable photo'). When data is changed since the operation started an error will be raised by the service routine.

For each service it is defined whether the consistent read option is supported or not. If not the service fails when it is asked to implement a consistent read.

IMPORTANT: When a service enforces a consistent read it only offers it correct when in the database the ROWDEPENDENCIES setting is activated for the OHI tables involved. This is a one time only database table reorganization action but will imply a large downtime to implement this for all tables.

3.4.1.3 Enforce_unchanged_since_scn

The setting is used typically in a change scenario to implement an optimistic locking algorithm. It should contain a numeric value which identifies an SCN value (System Change Number, a sequential change number assigned to changes and also to each committed transaction; please see the standard Oracle database documentation for more information).

If specified a non null value (and the called service supports it) the service should check for all records that will be changed (and perhaps in special cases also for other records which are retrieved in the operation, but this is service specific), whether they are not changed since the 'SCN moment' which is passed (a parameter value is passed for this). The SCN for the retrieved records may not be younger (larger) than the provided SCN (the third parameter specifies this).

Of course for updates/deletes an explicit lock with nowait is always implemented for the records that are affected (immediately before they are changed), disregarding this functionality. This to prevent the service 'hangs' on an outstanding lock. Internally the service determines the SCN for a record as part of the 'select for update' or after that select has succeeded (because the record is locked from that moment on, until the transaction ends or fails and rollbacks).

When a value zero is passed the 'SCN moment' at the start of the service call will be used (this is a special situation).

IMPORTANT: The same remark as in previous paragraph regarding ROWDEPENDENCIES applies for a correct working of this functionality.

3.4.1.4 Combining these settings

B y combining the previous 2 settings it can be specified for a service call to check whether data is read consistent during a retrieve of this data and to check whether the data to be changed during a (slightly) later service call is not changed in the meantime. This is done by remembering the SCN call moment of the retrieve call (which is returned in the returning context which is described later) and passing it on to the change call.

For retrieve only functionality the last setting is normally not relevant. Typically only the consistent read option will be supported but their may be exceptions.

For service calls that change data it can be useful to specify both the enforce_consistent_read and the enforce_unchanged_since_scn settings. When both

options are supported this means the routine will support a consistent read for 'supporting data' that is retrieved, but for the data which will be changed the enforce_unchanged_since_scn value that is passed, is used to check whether it has not changed since that moment.

3.4.2 Return Context

The return context looks like shown below.

```
create or replace force type svl_return_context_tp
as object
(
  message_info       svl_message_info_tp
, scn_call_moment    number
, constructor function svl_return_context_tp
  return self as result
)
```

It typically can contain a list of regular error messages that occurred during the call. These are passed using the message_info type, the first setting.

Next to the message info also the database SCN number (in fact a 'moment in time identifier' of the last change in the system at a specific moment) of the moment the service call started is returned, if consistent_read is enforced (!) and supported. This can be used in subsequent calls when the calling environment wants to make sure that data which is accessed/changed in the later call is not changed since the call which returned the SCN (see previous paragraphs).

With using the combination of the return context and the call context in subsequent calls an optimistic locking approach can be implemented in interface applications that require this. This is especially useful because of the stateless behaviour of the service calls, it is not possible to use a locking strategy with actual (row) locks.

3.4.2.1 Message info

The definition of the message info is shown below.

```
create or replace force type svl_message_info_tp
as object
(
  list           svl_message_lst_tp
, top_severity   varchar2(1)
, error_stack    varchar2(32000)
, constructor function svl_message_info_tp
  return self as result
)
```

The first variable in this message info object contains the actual list of messages. When there is a list of messages present (one or more messages in the list) this means a (functional) error has occurred which is handled in the exception handler of the called service routine. The message info variable contains in such a situation also the most severe level of the messages in the message list (top_severity). When for example 2 informational messages and one error message are in the list the top_severity is 'Error' ('E').

In the situation that an exception has occurred which is handled and put in the message list also the plsql error stack can be passed. This can help in detecting the cause of programming or application problems.

A message on the message list is structured as shown below:

```
create or replace type svl_message_tp
as object
  ( message_code               varchar2(32)
  , severity                  varchar2(1)
  , severity_desc              varchar2(100)
  , message_text               varchar2(2000)
  , help_text                 varchar2(2000)
  , lang_code                 varchar2(3)
```

```
  , constructor function svl_message_tp
    return self as result
  )
```

This shows a code is passed that uniquely identifies the specific message type (of course more of the same messages can be in the list) next to the severity, the description of the severity, the actual text of the message (in the language as defined by the preference of the user with which identity the service is called, which language is passed back in the lang_code variable) and optionally a help text which applies to that message.

3.5 Error and exception handling

When a call fails it can fail gracefully or by throwing an exception. When it fails gracefully a list of messages is returned in the returning context.

In case of an ungraceful failure, which is not handled by an exception handler in the service, an exception is thrown. Within the database this is typically an Oracle exception.

At the web services side SOAP fault messages reflect both situations. The XSD below specifies for this purpose a functional and a technical fault.

```
 <xsd:element name="functionalFaultType">
 <xsd:annotation>
  <xsd:documentation>SOAP Fault which is returned when a functional regularly
handled error has occurred.</xsd:documentation>
 </xsd:annotation>
 <xsd:complexType>
  <xsd:sequence>
   <xsd:element minOccurs="0" maxOccurs="unbounded" name="messages"
               type="faultMessageType"/>
  </xsd:sequence>
 </xsd:complexType>
</xsd:element>
<xsd:element name="technicalFaultType">
 <xsd:annotation>
  <xsd:documentation> SOAP Fault which is returned when a technical unhandled error
has occurred. </xsd:documentation>
 </xsd:annotation>
 <xsd:complexType>
  <xsd:sequence>
   <xsd:element name="code" type="xsd:string" minOccurs="0"/>
   <xsd:element name="message" type="xsd:string"/>
  </xsd:sequence>
 </xsd:complexType>
</xsd:element>
<xsd:complexType name="faultMessageType">
 <xsd:sequence>
  <xsd:element name="severityText" type="xsd:string"/>
  <xsd:element name="severityCode" type="xsd:string"/>
  <xsd:element name="messageText" type="xsd:string"/>
  <xsd:element name="messageCode" type="xsd:string"/>
  <xsd:element name="helpText" type="xsd:string" minOccurs="0"/>
 </xsd:sequence>
</xsd:complexType>
```

All application raised messages or handled exceptions are returned as functional faults and can be handled as such. As long as functional faults are returned it is clear the service call itself is still being executed along all the layers in the technology stack but somewhere during processing in the application functional errors or a handled exception occurs.

When an unhandled exception occurs this is typically returned as a technical fault and can be handled differently. This is more severe and can have all kind of causes. For example the database can be unreachable or down. But it cannot be said by definition that processing should be aborted, that is dependent on the cause of the technical fault. It might well be that the next call succeeds again because there only was a switch over from for example a failing network component to a replacement.

Beware that the structure of the XML messages, both the request as well as the response message, will be validated against the WSDL/XSD definitions. Errors will be returned when the message does not comply with the definition.

3.6 Transaction handling

For web services that can change data (so which implement more than only 'retrieve' operations) the change operation will commit automatically or fail in a consistent way (rolling back partially executed transactions). This behaviour is implemented automatically in the web service implementation.

When using the plsql implementation the calling code is responsible for committing or rolling back partially executed operations. When exceptions are handled in an exception handler be sure a rollback is executed in the exception handler to prevent partially executed transactions are committed (although this will never result in inconsistent data, that is guarded by the business rule implementation layer in the database), resulting in potentially unwanted changes.

3.7 Differences between provider plsql and web service 'service layer implementation'

The plsql and web service 'services' are very similar. In fact the plsql services are 'wrapped' into a web service. Typically there is one database plsql package supporting the operations of the corresponding web service.

In the service layer implementation the supporting web service plsql packages are named SVL_WS_<service> and contain the operations as packaged procedures. A standard function is_alive is offered that is wrapped as web service to check the complete web service technology stack is working fine. The is_alive function returns the version number of the associated package.

The operations, implemented as packaged procedures, accept as input and output parameters of user defined object types that are similar formatted as the input and response messages in the corresponding WSDL/XSD definitions.

In the current release it is not possible to access the call context and the return context in the web service implementation, as they are not present in the WSDL definition. In a future release this will be enhanced (at this moment they are of no use because the current services do not support the consistent read and optimistic locking functionality).

There is one exception: the calling user name can and must be specified in the Back Office properties file for each web service. Please read the installation and configuration manual for how to specify this calling user name. It will be passed as the user context part for the call context as described earlier.

So to state it more clearly, in the plsql implementation it is possible to specify and access the call and return context directly where this is not (yet) possible in the web service implementation.

3.8 Example plsql usage scenario

When the plsql implementation is used there are lots of possibilities in how to use the services. They can be combined with SQL and plsql to retrieve and change data directly or they can be used solely.

Below a very simple example is given how to retrieve some policy data using a plsql program:

```
declare
  l_call_context    svl_call_context_tp :=
svl_call_context_tp(svl_user_context_tp('MANAGER'),svl_yes_no_tp('N'),svl_yes_no_tp('N'),0);
  l_return_context   svl_return_context_tp;
  l_pol_details_tp   svl_policy_and_details_tp := svl_policy_and_details_tp();
begin
  svl_ws_policy_pck.get_pol_detail_by_pol_num_int
  ( pi_pol_nr_tp     => svl_policy_number_internal_tp(17553)
  , pi_call_context   => l_call_context
  , po_pol_detail_tp  => l_pol_details_tp
  , po_return_context => l_return_context
  );
  if l_return_context.message_info.messages.count > 0
  then
    dbms_output.put_line(l_return_context.message_info.error_stack);
    for i in 1..l_return_context.message_info.messages.count loop
      dbms_output.put_line('Msg('||i||'): '||l_return_context.message_info.messages(i).message_text);
    end loop;
  else
    dbms_output.put_line('No errors occurred.');
    for i in 1..l_pol_details_tp.membership_lst.count loop
      dbms_output.put_line('Mmb('||i||'): since=<'||l_pol_details_tp.membership_lst(i).start_date||
      '> name=<'||l_pol_details_tp.membership_lst(i).member_tp.formatted_name.formatted_name||
      '> sofi=<'||
      l_pol_details_tp.membership_lst(i).member_tp.social_security_number.social_security_number||'>');
    end loop;
  end if;
end;
```

As you can see the service operation SVL_WS_POLICY_PCK.GET_POL_DETAIL_BY_POL_NUM_INT is used. As calling user the known application username MANAGER is used. This code is executed using a database account created for this purpose (as using the application object owner directly is not supported).

3.9 Example provider web service usage scenario

In the situation of a web service of course the WSDL URL should be used to access the WSDL. The system administrators who deploy the web services should provide this URL.

A typical WSDL URL could be:

http://<servername>:<port>/OHIBOWebservices/OhiPolicyService?wsdl

To test whether a service is technically working a tool like soapUI can be used which creates requests for the different operations.

The isAlive operation can be used for the technical test.

However, a simple example, which is similar to the plsql example in the previous paragraph, can also be used:

```
<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
xmlns:v1="http://www.oracle.com/insurance/ohibo/policy/policymessages/v1">
   <soapenv:Header/>
   <soapenv:Body>
      <v1:getCurrentPolicyDetailsByPolicyNumberInternalRequestType
policyNumberInternal="17553"/>
   </soapenv:Body>
</soapenv:Envelope>
```

This will return a SOAP message containing the contents of the returned policy message structure or a SOAP fault message structure in case of problems.

4 Consumer web services

Starting with release 2012.01 a first implementation is offered of consumer services that are consumed within the database. These services support existing batch functionality that consumes services in the outside world. Where outside world is defined as 'outside of the OHI application'.

For implementing these services it is expected that for amongst others security and traceability reason an internal 'facility' can be called that provides these services. This intermediate 'gateway' (typically a proxy or a service bus) implements the call to the real outside world. In this way for example standardized security solutions can be used to protect the communication to the outside world, independent from the OHI Back Office application.

This chapter focuses on implementation aspects for the consumer services.

4.1 Technical principles

The following technical principles are followed and may be of influence for how you realise your code to access the Service Layer consumer web services.

1. SOAP 1.1 is used.

   Rationale: this is by far still the most widely used common standard and supported by almost all web service toolkit implementations.

2. The OHI provided WSDL's are currently look-a-likes for the outside world WSDL's that are offered by VECOZO.

4.2 Gateway setup

Internally a proxy or a Service Bus should be setup to provide the provider web services functionality that can be consumed by the consumer web services.

To know the functionality that should be implemented a WSDL is delivered per consumer web service that describes the interface to be offered.

These OHI provided WSDL's are currently look-a-likes for the outside world WSDL's that are offered by Vecozo. It should be quite easy to identify how to map the data of the external and the 'OHI' WSDL.

For more information currently the Functional Specification of theme M-2647 should be used.

4.3 Configuration

Currently only Dutch localisation consumer web services exist. The table below shows which Back Office parameter settings should be configured for each service. The values should identify the internal gateway you configure.

Code	Name	Back Office parameters

Code	Name	Back Office parameters
FSH1009S	Uitvoering fraudecontrole (VECOZO)	1. EVREndPoint 2. EVRProxyHost 3. EVRProxyPort
FIN2114S	Aanmaken en versturen borderel ambtshalve verzekerden	1. AmbtshalveEndPoint 2. AmbtshalveProxyHost 3. AmbtshalveProxyPort
ZRG1293S	Controle op premieachterstand VECOZO	1. PremieAchtEndPoint 2. PremieAchtProxyHost 3. PremieAchtProxyPort
ZRG1298S	Opzegservice VECOZO	1. OpzegSrvceEndPoint 2. OpzegSrvceProxyHost 3. OpzegSrvceProxyPort
ZRG2221S	Aanmaken en versturen AVG-bestand	1. AVGEndPoint 2. AVGProxyHost 3. AVGProxyPort
ZRG3078S	Genereren machtiging retourbericht (XML)	1. MachtigingEndPoint 2. MachtigingProxyHost 3. MachtigingProxyPort

4.4 Error and exception handling

When a consumer web service call fails it fails by throwing an exception. Because the calls are implemented inside the database this is typically an Oracle exception with the error code ORA-29532. These errors are stored as messages that occurred during the batch that executed the consumer services.

The ORA-29532 error code indicates a java exception; the actual error is shown in the error message that follows after the error code.

Below is a non-exhaustive list of possible errors that can occur when a consumer service is called by the OHI Back Office application:

Error	Cause
ORA-29532: Java call terminated by uncaught Java exception: java.rmi.RemoteException: java.rmi.RemoteException:; nested exception is: HTTP transport error: javax.xml.soap.SOAPException: java.security.PrivilegedActionException: **javax.xml.soap.SOAPException: Message send failed: Connection refused**	No web server available at the given location
ORA-29532: Java call terminated by uncaught Java exception: java.rmi.RemoteException: java.rmi.RemoteException:; nested exception is: HTTP transport error: javax.xml.soap.SOAPException: java.security.PrivilegedActionException: **oracle.j2ee.ws.saaj.ContentTypeException: Not a valid SOAP Content-Type: text/html; charset=iso-8859-1**	A web server is available at the given location, but does not accept SOAP messages or cannot respond to the requested message (unknown request)
ORA-29532: Java call terminated by uncaught Java exception: java.rmi.RemoteException: **oracle.j2ee.ws.common.encoding.DeserializationException:deserialization error: java.lang.IllegalArgumentException**	Unknown or invalid (response) message: invalid value
ORA-29532: Java call terminated by uncaught Java exception: java.rmi.RemoteException: **java.lang.NullPointerException:null**	Unknown or invalid (response) message: invalid namespace

Error	Cause
ORA-29532: Java call terminated by uncaught Java exception: java.rmi.RemoteException: **java.rmi.RemoteException:Error parsing envelope: (1, 1) Start of root element expected.; nested exception is: javax.xml.soap.SOAPException: Error parsing envelope: (1, 1) Start of root element expected.**	Unknown or invalid (response) message: empty message
ORA-29532: Java call terminated by uncaught Java exception: java.rmi.RemoteException: **oracle.j2ee.ws.common.encoding.DeserializationException:unexpected element name: expected={urn:http://www.oracle.com/insurance/ohibo/SVL1001C:messages:isevr:v1}EvrStatus, actual={urn:http://www.oracle.com/insurance/ohibo/SVL1001C:messages:isevr:v1}EvrStatuss**	Unknown or invalid (response) message: Wrong name of element

5 Appendix A – Provider web services documentation per service

Please use the WSDL that can be retrieved when a web service is deployed.

When services are changed the functional specification contains the latest changes.

Currently the provider web services offer no consistent read or locking functionality and can only retrieve data through a number of operations, except for services that implement changes.

The SVL_WS% packages can be used to get an overview of the existing services and their operations (the packaged procedures).

For more information please see the Business Function that start with WS_ within application system SVL as defined in the Designer Repository. The operations are documented as SubFunctions for the Business Services defined (6 at this moment).

Each operation is documented through either HTML in the Description field (be sure you use an HTML editor to open this field) or is documented in a stored File (also present as 'File' in the application system SVL) that is named in the Description field.

Attention: Starting with release 2011.03.2 the PreAuthorization provider web service WSDL has been extended with an additional 'Herkomst' field. For the corresponding consumer web service the official Vecozo WSDL has been extended with an optional field that is not supported by Vecozo.

When the provider web service call specifies a value for the 'Herkomst' field the consumer call will fill the non supported Vecozo field. This means the call will not be accepted by Vecozo.

This functionality is introduced for the situation where an internal 'intervening component' (typical a middleware solution like a service bus) is used as intermediate and there is a requirement to distinguish between sources for pre authorization requests. This can be helpful in using the pre authorization web service for more source than the Vecozo pre authorization web service.