BEAAquaLogic
Data Services
Platform™

Release Notes

Contents

Release Notes

Release Notes

BEA AquaLogic Data Services Platform Version 2.0.1

Latest Revision Date: November 18, 2005

*** NEW *** DSP Certified for WebLogic Platform 8.1 SP5

The BEA AquaLogic Data Services Platform 2.0.1 now supports WebLogic Platform 8.1 Service Pack 5 (SP5). No upgrade is needed. However, if you are planning on using DSP 2.0.1 with SP5 see release note CR247416 related to interdomain transactions between WebLogic Server 8.1 SP4 and 8.5 SP5.

About BEA AquaLogic Data Services Platform

Data Services Platform (DSP) provides on-demand access to, and aggregation of, information in relational databases, Web services, Java functions, XML files, delimited files, and other types of disparate data. During development, DSP is accessed as an application running under WebLogic Workshop. This is where you develop your data services, including query and relationship functions.

Once developed, application developers — using access technologies such as the Data Services Platform mediator API, JDBC, SQL, or Data Services Platform controls — can invoke DSP query functions as a means of providing their applications with access to integrated, updateable data from their enterprise.

Metadata, security, and cache management facilities are provided through the Data Services Platform Console (ldconsole).

Note: Data Services Platform was formerly named Liquid Data. Some artifacts of the original name remain in the product, path, and components.

These Release Notes are subject to revision between releases. The most recent version can always be found at:

> http://edocs.bea.com/aldsp/docs20/relnotes/index.html

Items modified or added are so marked by date.

This document includes the following topics:

- About Data Services Platform 2.0.1
- Product Limitations and Workarounds
- References

About Data Services Platform 2.0.1

Data Services Platform 2.0.1 is a service pack release primarily designed to validate product durability and address known issues. This has included:

- Rebranding of Liquid Data to its new name and version number: BEA AquaLogic Data Services Platform 2.0. This included numerous changes to documentation and to the user interface. As noted above, some artifacts of the original name remain in the product, installation path, and components remain.

- Incorporation of a number of product patches developed for various specific installations.

- Performance improvements.

Product Limitations and Workarounds

Table 1 lists known, potentially encountered limitations associated with the current BEA AquaLogic Data Services Platform 2.0.1 release. Information regarding these limitations includes a CR (Change

Request) number for each problem, applicable platform, a detailed description of the problem and workarounds, where applicable.

Please contact BEA customer support at http://support.bea.com for assistance in tracking unresolved problems.

Table 1 Known Product Limitations and Workarounds

Topic		Details
CR242938	**(11/05)**	**Interdomain transactions between WebLogic Server 8.1 SP4 and 8.1 SP5 require a command line flag to be passed.**
Description		Running a Web service created for a Data Service control from a remote machine gives the following exception: Current server is the coordinator and transaction is not found.
Platform		All.

Workaround		When using interdomain transactions between 8.1 SP4 and 83.1 SP5 a command-line flag to the 8.1 SP5 domain is needed for DSP if the Data Service control and the calling application are in separate domains with one running 8.1 SP4 and the other 8.1 SP5.
		
		`-Dweblogic.transaction.SecurityInteropMode=compatibility`
		More details on the JTA transaction are available at:
		
		`http://e-docs.bea.com/wls/docs81/ConsoleHelp/jta.html#1106135`

CR242938	**(9/05)**	**Multi-dimension soap arrays are not supported in RPC mode.**
Description		Currently the Web services wrapper provided by DSP only supports single-dimension arrays in RPC style Web services
Platform		All.
Workaround		None.
CR224815	**(9/05)**	**The initial invocation of a Web service call from the application server typically takes more time than subsequent calls. If the timeout value is less than the time required for the first call, the alternate expression (typically a timeout error) will be evaluated.**

Topic	Details
Description	There is "startup overhead" the first time that a web service is invoked. The overhead can exceed the timeout threshold, leading to the specified timeout error.
Platform	All.
Workaround	When setting timeout on expressions that have a Web service invocation, set the timeout value to be greater than the measured amount of time required for the first invocation.
CR238736 (9/05)	**Wildcards may not work when specifying a catalog name.**
Description	The JDBC driver may not support the use of wildcards (i.e., WIRE% for WIRELESS) as a means of identifying database catalogs. The result is that no catalog is found when a wildcard is used to specify an available database catalog.
Platform	All.
Workaround	Use the exact name, not wildcards, when specifying a catalog name.
CR239369 (9/05)	**XQueries may generate invalid SQL for databases not supporting UPPER and LOWER (SQL-92).** **Also, empty input handling for base databases (databases not specifically support) as well as Oracle deviates from the XQuery specification when UPPER(null) or LOWER(null) is pushed down to the database level.**
Description	There are two aspects to this problem: 1. XQueries containing upper-case() or lower-case() functions are pushed down for database processing as UPPER and LOWER. Some databases may not support these SQL-92 keywords, however. In such cases the generated SQL will be invalid and will upon execution fail. 2. Similarly, input handling by base databases (as well as Oracle databases) may not match the XQuery specification. The reason for this is that while the XQuery specification requires that functions *return an empty string* if input is an empty sequence. However, when these functions are pushed down, they *return an empty sequence* instead. This happens because LOWER(NULL) is NULL in SQL.
Platform	All platforms running base databases and Oracle databases

Topic		Details
Workaround		Use the fn-bea:Fence() function to prevent pushdown of upper-case() or lower-case() functions to the database. Example: `lower-case(fn-bea:fence(...))`
CR237186	(9/05)	**SDO update sequence is not preserved.**
Description		SDO update does not preserve the sequence of update objects; instead it updates in alphabetical order (example: **a**ddress, **c**redit card, **c**ustomer). Since the insertion order is automatic, the update will fail unless the update elements match alphabetical order.
Platform		All.
Workaround		If the dependency order matches alphabetical order, the update will be successful.
CR238044	(9/05)	**Metadata search using the '?' wildcard may not return expected results.**
Description		In metadata browser search the "?" wildcard should match with exactly one valid character of any value. However, the "?" wildcard actually matches any exact number of characters or a subset. For example, searching for "CUSTOMER_???" will match both "CUSTOMER_ABC" and "COSTUMER_AB". While this is in keeping with certain Java wildcard behavior, it is inconsistent with standard Windows search behavior.
Platform		All.
Workaround		None.
CR207637 (9/05)		**An exception appears for XQuery functions accessing metadata derived from Microsoft SQL Server stored procedures containing xs:decimal.**
Description		When importing a stored procedure from Microsoft SQL Server, the BEA JDBC driver incorrectly maps SQL decimal type to schema integer (xs:int) type.
Platforms		All.
Workaround		During stored procedure import, change the data type from xs:int to xs:decimal. Alternatively, you can change the imported data service's metadata to specify the schema type for the affected column to be xs:decimal.
CR203394		ROWTYPE input cursor is not supported when creating a data service from a stored procedure.

Topic	Details
Description	Stored procedure IN and INOUT cursors containing ROWIDs are not currently supported for metadata import.
Platform	All.
Workaround	Avoid importing metadata on stored procedures which required use of IN or INOUT ROWID parameters.
CR218040	**The query engine cannot retrieve data from a Sybase data source when the generated SQL contains functions which have dependencies on Java services.**

Description	When using Sybase relational data as a data source, some SQL functions such as LTRIM and RTRIM work correctly in the database only when Java services are enabled by running the following in the SA role.
	
	`sp_configure, "enable_java", 1`
	
	The error message indicating this symptom is:
	
	`Cannot run this command because Java services are not enabled. A user with System Administrator (SA) role must reconfigure the system to enable Java.`

Platform	All platforms running Sybase.
Workaround	Enable Java services while in the SA role. Refer to Sybase documentation for details.
CR210000	**The XQuery 1.0 processing-instruction() node test is not fully supported in the BEA XQuery processor.**
Description	Using processing-instruction("name") ignores the "name" and instead retrieves all processing-instruction nodes.
Platform	All.
Workaround	None available
CR221145	**WSDLs with multiple services are not supported.**
Description	A limitation of the WebLogic web services stack interferes with the ability to handle WSDL files with multiple data services.
Platform	All.

Topic	Details
Workaround	Split multiple services into multiple WSDL files.
CR214585	**Erroneous results may occur when using fn:matches() with a regular expression containing a caret (^).**
Description	The match beginning-of-line operator (^) in regular expressions produces erroneous results when used with `fn:matches()`.
Platform	All.
Workaround	None available.
CR215251	**Identifiers within two characters of the maximum length allowed by the database may result in an error.**
Description	Some databases place limits on the length of identifiers (30 in the case of Sybase). Data Services Platform places single quotes around pushdown queries, effectively reducing the maximum identifier length by two characters (28 in the case of Sybase).
Platform	All platforms running Sybase (and possibly other) databases.
Workaround	Possible options include renaming the table or creating a view with a shorter name.
CR203174	**Some data sources may not appear when attempting to import metadata.**
Description	The functionality to manage JDBC data sources and connection pools during metadata import is limited to the same level of functionality provided by BEA Workshop (Tools →WebLogic Server →DataSource Viewer).
Platform	All.
Workaround	For full functionality use the WebLogic Administration Console to manage JDBC data sources and connection pools.
CR209659	**SOAP 1.2 for Web services runtime is not supported.**
Description	This version of Data Services Platform cannot invoke SOAP 1.2 based Web services.
Platform	All.

Topic	Details
Workaround	None available.
CR222822	**Linux application built in Workshop may fail with the message "Error creating temporary file".**
Description	While building an application inside Workshop on Linux, you may get the above error when 1) the application includes a large number of schema files, and 2) when the Platform Installation is under a different user's name than the user running Workshop.
Platform	Linux.
Workaround	Ensure that the Workshop instance is running under the same name as was used for Platform installation.
CR226019	**Access control policies associated with a data service function may disappear if the function's number of parameters is changed.**
Description	A data service function's signature is it's QName and the number of parameters (arity) of the function. If you set security policies on a function and changes the number of parameters to the function, then this function is treated as a new function and the previously set policies will no longer be in effect.
Platform	All
Workaround	If a functions arity is changed, the user will have to reapply security policies to that function.
CR227486	**The BEA JDBC XA driver for DB2 returns the error XAER_RMERR when a local transaction read is followed by a global transaction read.**
Description	Pertains to a known JDBC driver problem described in BEA CR229071.
Platform	All
Workaround	For the 3.4 JDBC driver, set the driver property to: `AllowImplicitResultSetCloseForXA=false`
CR229758	**Data Services Platform Web service generation may fail to compile if the original Web service was created in Workshop.**

Topic	Details
Description	If you have a data service created from a Workshop-originated Web service and then, subsequently, you turn your data service into a web service via a Data Services Platform control, the project build will fail due to a duplicate schema element error.
Platform	All

Workaround	Assume that you have a Workshop-generated Web service named getCustomerOrderByCustomerID. A WSDL for this Web service with default targetnamespace would create the following schema elements:
	
	`{http://www.openuri.org/}getCustomerOrderByCustomerI`
	`D`
	and
	`{http://www.openuri.org/}getCustomerOrderByCustomerI`
	`DResponse`
	
	When you import this WSDL, the same schema elements will be in the imported schema file and will be associated with the data service read function. When you build a Data Service control for this read function and generate a Web service from that control, a Web service function is created:
	
	`org.openuri.GetCustomerOrderByCustomerIDResponseDocu`
	`ment`
	`getCustomerOrderByCustomerID(org.openuri.GetCustomer`
	`OrderByCustomerIDDocument p0)`
	
	When compiling this Web service file, now workshop generates two identical schema elements each for

Topic	Details
	When compiling this Web service file, Workshop now generates two identical schema elements each for:
	`{http://www.openuri.org/}getCustomerOrderByCustomerI` `D`
	and
	`{http://www.openuri.org/}getCustomerOrderByCustomerI` `DResponse.`
	Thus the schema elements clash.
	If you modify either the target namespace of the final JWS or change the function name to a different one such as:

	`org.openuri.GetCustomerOrderByCustomerIDResponseDocu`
	`ment`
	`getCustomerOrderByCustomerIDNEW(org.openuri.GetCusto`
	`merOrderByCustomerIDDocument p0)`

	Then the types generated will be as follows and there will be no clash.
	`{http://www.openuri.org/}getCustomerOrderByCustomerI` `D`
	`{http://www.openuri.org/}getCustomerOrderByCustomerI` `DResponse`
	`{http://www.openuri.org/}getCustomerOrderByCustomerI` `DNEW and`
	`{http://www.openuri.org/}getCustomerOrderByCustomerI` `DNEWResponse`
	Thus the schema clash will be avoided.
CR213916	**BEA Informix JDBC driver does not return nullability information.**
Description	The BEA Informix driver does not return information about table column nullability (that is, it is marked as unknown). Thus during metadata import the minOccurs of the elements corresponding to the columns in the generated XML schemas is set to 0.
Platform	All platforms running Informix.
Workaround	Modify the imported metadata files by changing the minoccurs value for the nullable columns from 0 to 1.

Topic	Details
CR214983, CR211701, CR201821	**MSSQL VARIANT datatype has only limited support.**
Description	There are two limitations with this MSSQL VARIANT data type (sql_variant): 1. For SQL_VARIANT data type update will fail. 2. You cannot read a null value for the SQL_VARIANT data type.
Platform	All
Workaround	None available
CR218449	**Data services cannot have multibyte namespace prefixes.**
Description	During metadata import, tables with multibyte names may create multibyte namespace prefixes based on the table name's first few characters. The resulting data service file will not be parsed correctly because the XMLBeans parser is unable to handle multibyte namespace prefix value.
Platform	All
Workaround	None available
CR211377	**In Source View syntax highlighting may become off by a character if certain sequences appear in the query.**
Description	If your query has any of the following patterns: • Numeric XML character entity references, e.g., &#20 • Alphabetic XML character entity references, e.g., &amp, &apos, &gt, &lt, and &quot • The {{ and }} tokens. • Escaped single and double quotes. then syntax highlight may become off by one.
Platform	**All**
Workaround	**None available**
CR219009	**Data Services Platform IDE Test View is limited to a flat structure when performing a SDO update.**

Topic	Details
Description	Through Data Services Platform Test View Data you can accomplish typical read and update actions. For example, you can create new SDO records as part of update steps. However, when adding a new SDO record that has a nested structure, the template generated will only have first-level (i.e., no nested) elements.
Platform	All
Workaround	You can add subsequent nested SDO records by right-mouse clicking on the container node.
CR221015	**During metadata import the BEA Sybase JDBC driver may not display all tables to which user has authorized access.**
Description	During metadata import the BEA Sybase JDBC driver may not show all tables which have been granted access to the user.
Platform	All platforms accessing Sybase through the BEA Sybase JDBC driver.
Workaround	This is a BEA Sybase driver limitation. For the import purpose, you can change to dbo user to see the full complement of available tables.
CR221956	**Repeated builds of a Workshop application containing one or more Web service projects can run out of memory.**
Description	If you have a Workshop application that has a large number of Web service files (.jws) in a Web service project, the build of such an application can cause a memory leak and, after certain number of repeated attempts, Workshop may run out of memory. This, in turn, can lead to various class loading errors.
Platform	All
Workaround	Obtain and apply the Workshop patch associated with CR226814. This patch corrects the memory leak problem described above.
CR202963	**When using BEA Oracle JDBC driver with a TIMESTAMP values, stored procedures are truncated.**
Description	When using the BEA's Oracle JDBC driver, if a stored procedure returns a TIMESTAMP value then the value gets truncated at the milliseconds level. For example, if the value was 1997-01-31 09:26:50.124 then the stored procedure will return a 1997-01-31 09:26:50.0 value.

Topic	Details
Platforms	All platforms running Oracle with the BEA Oracle JDBC driver.
Workaround	Use the Oracle JDBC driver stored procedures that returning TIMESTAMP values.
CR223429, CR228802	**Sybase JDBC driver does not support a getBlob() call.**
Description	The Data Services Platform cache configuration does not work if using Sybase JDBC driver because the configuration implementation uses a getBlob() call on the JDBC driver. The Sybase JDBC driver does not support getBlob().
Platforms	All platforms running Sybase with the Sybase JDBC driver.
Workaround	Use the BEA JDBC driver for Sybase databases when utilizing Sybase as the Data Services Platform cache data source.
CR214730	**SQL Server JDBC driver incorrectly renders the tinyint maximum value.**
Description	The SQL Server tinyint maximum value of 255 gets interpreted as -1 by the Microsoft SQL Server JDBC driver.
Platforms	All platforms running SQL Server with the SQL Server JDBC driver.
Workaround	Use the BEA JDBC driver for SQL Server.
CR223486, CR226171, CR226239	**The Informix JDBC driver does not support standard JDBC syntax for specifying TIMESTAMP values.**
Description	The Informix native driver doesn't support standard JDBC syntax for specifying TIMESTAMP values. For example: 1979-03-01 00:00:00.0 is not supported.
Platforms	All platforms running Informix with the Informix JDBC driver.
Workaround	Use the BEA JDBC driver for Informix.
CR199675	**The BEA JDBC driver for Oracle does not support UROWID column type for data retrieval.**
Description	When using BEA JDBC driver for Oracle, retrieving UROWID returns an error, identified by the following message: `[BEA][Oracle JDBC Driver]Internal error: Net8 protocol error`

Topic	Details
Platforms	All platforms running Oracle with the BEA JDBC driver.
Workaround	Use the Oracle JDBC driver if your data contains UROWID column type.
CR212515	**The Oracle stored procedure returning PL/SQL RECORD, BOOLEAN, or table with non-scalar element types is not supported.**
Description	Oracle stored procedure limitations are detailed in the following currently-available document: http://www.stanford.edu/dept/itss/docs/oracle/9i/java.920/a96654/ref.htm#1007714
Platforms	All platforms using Oracle.
Workaround	None available
CR202041	**Metadata for SQL Server stored procedures returning CURSOR output cannot be created.**
Description	When importing metadata from SQL Server, stored procedures that return CURSOR output are not supported.
Platforms	All platforms using SQL Server.
Workaround	Modify the imported data service file to the identify the correct data type for the cursor.
CR227440	**Metadata for DB2 stored procedures returning CLOB data cannot be created.**
Description	When importing metadata from DB2, stored procedures returning CLOB data are not supported.
Platforms	All platforms using DB2.
Workaround	None available.
CR202962	**Oracle stored procedures containing CHAR or NCHAR as input can cause a Server error.**
Platform	All

Topic	Details

Description	If you have Oracle stored procedures that use an INOUT parameter, you may get an error when you run a stored procedure using Data Services Platform. The error appears as:
	
	`java.lang.RuntimeException: ORA-01460: unimplemented or unreasonable conversion requested`
	
	`ORA-06512: at "WIRELESS.SP_CHAR", line 17`

Workaround	Modify your stored procedure call by reducing the size of the INOUT parameter using TRIM. See Listing 1, below.

Listing 1   Sample code related to CR202962

```
CREATE OR REPLACE PROCEDURE WIRELESS.SP_CHAR
(P_CHAR_IN IN CHAR,
P_CHAR_OUT OUT CHAR,
P_CHAR_INOUT IN OUT CHAR,
P_ID_OUT OUT VARCHAR2 )
IS
TEMP VARCHAR2(10);
BEGIN
SELECT C_ID INTO P_ID_OUT
FROM WIRELESS.ALL_DATATYPES
WHERE C_CHAR = P_CHAR_IN;

SELECT C_CHAR INTO P_CHAR_OUT
FROM WIRELESS.ALL_DATATYPES
WHERE C_ID = '2';

SELECT C_ID INTO TEMP
FROM WIRELESS.ALL_DATATYPES
WHERE C_CHAR = P_CHAR_INOUT;

SELECT 'WORK' INTO P_CHAR_INOUT
FROM WIRELESS.ALL_DATATYPES
WHERE C_ID = TEMP;
END;
/
```

to adjust the size of PCHAR_INOUT using TRIM (see highlighted code)

```
CREATE OR REPLACE PROCEDURE WIRELESS.SP_CHAR
(P_CHAR_IN IN CHAR,
P_CHAR_OUT OUT CHAR,
```

```
P_CHAR_INOUT IN OUT CHAR,
P_ID_OUT OUT VARCHAR2 )
IS
TEMP VARCHAR2(10);
ACHAR CHAR(500);

BEGIN
ACHAR := trim(P_CHAR_INOUT);

SELECT C_ID INTO P_ID_OUT
FROM WIRELESS.ALL_DATATYPES
WHERE C_CHAR = P_CHAR_IN;

SELECT C_CHAR INTO P_CHAR_OUT
FROM WIRELESS.ALL_DATATYPES
WHERE C_ID = '2';

SELECT C_ID INTO TEMP
FROM WIRELESS.ALL_DATATYPES
WHERE C_CHAR = ACHAR;  // used to fail here

SELECT 'WORK' INTO P_CHAR_INOUT
FROM WIRELESS.ALL_DATATYPES
WHERE C_ID = TEMP;

END;
/
```

References

Documentation for this product is available on an Online Documentation CD-ROM that ships with the product, and is also available on the Web at the BEA e-docs documentation site at the following location:

http://e-docs.bea.com/aldsp/docs20/index.html

To access the Data Services Platform documentation home page in a Web browser:

1. Click the PDF files button.

2. Select the document that you want to view or print.

Alternatively, from within a HTML document, click the View as PDF link on the top navigation bar to open a PDF version of the currently displayed document.

Information about BEA products including Data Services Platform can be found at:

http:dev2dev.bea.com

Documentation for all BEA products in both PDF and HTML format can be found at:

http://edocs.bea.com

If you do not have Adobe Acrobat Reader, you can obtain it from the Adobe Web site at:

http://www.adobe.com