


Extremely pureXML in DB2 10 for z/OS

Develop Java and COBOL applications
accessing SQL and XML data

Administer your XML and SQL data

Choose the best options for
installation and use



Paolo Bruni
Neale Armstrong
Ravi Kumar
Kirsten Ann Larsen
Tink Tysor
Hao Zhang



International Technical Support Organization

Extremely pureXML in DB2 10 for z/OS

January 2011

Note: Before using this information and the product it supports, read the information in “Notices” on page xvii.

First Edition (January 2011)

This edition applies to Version 10.1 of IBM DB2 for z/OS (program number 5605-DB2).

This document created or updated on January 9, 2011.

Note: This book is based on a pre-GA version of a program and may not apply when the program becomes generally available. We recommend that you consult the program documentation or follow-on versions of this IBM Redbooks publication for more current information.

Contents

Figures	ix
Tables	xi
Examples	xiii
Notices	xvii
Trademarks	xviii
Preface	xix
The team who wrote this book	xix
Now you can become a published author, too!	xx
Comments welcome.	xxi
Stay connected to IBM Redbooks	xxi
Chapter 1. Introduction	1
1.1 Importance of XML data	2
1.1.1 Growth of XML	2
1.1.2 The value of XML data	3
1.2 XML introduction	4
1.2.1 XML definitions	4
1.2.2 Document validity and well-formedness	5
1.2.3 XML Schema	9
1.2.4 Extensible Stylesheet Language	11
1.2.5 XPath	13
1.2.6 XQuery	15
1.2.7 XHTML	15
1.2.8 XSL, XSLT, Xpath, and XHTML examples	17
1.3 What is in this book	19
Chapter 2. XML and DB2 for z/OS	21
2.1 XML capabilities provided by DB2	22
2.1.1 Native XML data type	23
2.1.2 SQL/XML language	24
2.1.3 Hybrid data access	31
2.1.4 XML update	32
2.1.5 XML indexes	34
2.1.6 XML schema repository and schema validation	36
2.2 Supporting infrastructure	38
2.2.1 XSR installation steps	38
2.2.2 XSR installation validation	39
2.2.3 XSR setup troubleshooting	40
2.2.4 z/OS XML system services	42
2.3 Choice of tools	43
2.3.1 3270 based tools	43
2.3.2 GUI based tools	43
Chapter 3. Application scenario	45
3.1 Requirement for XML event logging and auditing	46
3.2 Application scenario	46

3.2.1	Using ISO 20022 with DB2 pureXML	47
3.3	Application code samples	48
3.3.1	DB2 SQL/XML programming pureXML	49
3.3.2	Using Java with DB2 pureXML	50
3.3.3	Using COBOL with DB2 pureXML	50
Chapter 4.	Creating and adding XML data	51
4.1	Creation of tables with XML columns	52
4.2	Storage structure for XML data	52
4.3	Multi-versioning concurrency control for XML	58
4.3.1	Example of improved concurrency with XML versions	58
4.3.2	Example of improved storage usage with XML versions	59
4.3.3	Storage structure for XML data with versions	60
4.4	Catalog queries to gather information	63
4.5	Display database command	66
4.6	Ingesting XML data	67
4.7	XML indexes	70
Chapter 5.	Validating XML data	73
5.1	XML schema validation	74
5.2	XML type modifier	74
5.3	Automatic validation	78
5.4	User-controlled validation	82
5.5	Determining whether an XML document has been validated	84
Chapter 6.	DB2 SQL/XML programming	87
6.1	Native SQL stored procedures and XML	88
6.1.1	Native SQL stored procedure example	89
6.1.2	XML error handling in native SQL procedures	91
6.1.3	Stored procedures development tools	93
6.2	Receiving XML messages from MQ	93
6.2.1	WebSphere MQ functions	94
6.2.2	DB2 stored procedure reading from MQ	95
6.2.3	DB2 MQ Listener automation	97
6.3	Audit queries (against logged XML messages)	99
6.3.1	Simple SQL/XML search examples	99
6.3.2	Choosing XML indexes	104
6.3.3	Verifying XML index usage	106
6.4	SQL/XML query techniques	106
6.4.1	Manipulating XML data with XPath functions	107
6.4.2	Filtering the rows returned with XMLEXISTS	108
6.4.3	Creating documents with publishing functions	109
6.4.4	Aggregating documents with XMLAGG	110
6.4.5	Enumerating all occurrences using XMLTABLE	111
6.4.6	Grouping data with XMLTABLE	112
6.5	User defined functions with XML	112
6.5.1	UDFs for reading from XML documents	112
6.5.2	UDFs for writing updates to XML documents	115
6.6	Triggers with XML	115
6.7	XML joins	116
6.7.1	XML to relational join	116
6.7.2	XML to XML join	117
6.8	XML with change data capture tools	118
6.8.1	Change data capture tools background	119

6.8.2 Using DB2 pureXML to receive CDC messages.	120
6.8.3 XML history objects.	125
Chapter 7. Using XML with Java	129
7.1 XML in Java	130
7.1.1 XML support in JDBC 3.0	131
7.1.2 XML support in JDBC 4.0	132
7.1.3 Constructing XML document in Java	133
7.1.4 Binary XML format in Java applications	135
7.2 The BankStmt application in Java	137
7.2.1 Setting up the environment	137
7.2.2 Insertion of rows with XML column values	139
7.2.3 Updates of XML columns	141
7.2.4 Retrieving XML data	144
7.2.5 Call stored procedure to shred XML	145
7.2.6 XSLT to transform XML document	147
7.2.7 Java interface to MQ.	151
Chapter 8. Using XML with COBOL	155
8.1 XML representation in COBOL	156
8.1.1 XML host variables in COBOL	156
8.1.2 Using non-XML variables for XML data	159
8.1.3 Using file reference variables for efficient insert and retrieval.	160
8.2 The BankStmt application in COBOL	161
8.2.1 Setting up the environment	162
8.2.2 Inserting XML documents	164
8.2.3 Updating XML documents.	167
8.2.4 Querying XML documents	171
8.2.5 Designing indexes.	172
8.2.6 Schema evolution	173
8.3 COBOL functions for manipulating XML	176
8.3.1 Generation of XML documents in COBOL	176
8.3.2 Shredding XML documents in COBOL	178
8.3.3 Validation of XML documents in COBOL	179
Chapter 9. Utilities with XML	181
9.1 CHECK DATA	182
9.2 CHECK INDEX	186
9.3 COPY	188
9.4 COPYTOCOPY	191
9.5 EXEC SQL	192
9.6 LISTDEF	193
9.7 LOAD	195
9.8 MERGECOPY	202
9.9 QUIESCE	203
9.10 REBUILD INDEX.	205
9.11 RECOVER INDEX and RECOVER TABLESPACE	206
9.12 REORG INDEX and REORG TABLESPACE	211
9.13 REPAIR.	215
9.14 REPORT	215
9.15 RUNSTATS.	219
9.16 UNLOAD	220
9.17 DSNTIAUL	225
9.18 DSN1COPY	228

Chapter 10. XML-related tasks for the DBA	229
10.1 Tasks regarding system setup	230
10.1.1 Setting up the XSR	230
10.1.2 Buffer pool for XML	231
10.1.3 Sizing XMLVALA and XMLVALS	231
10.1.4 Be up to date with maintenance	231
10.2 Tasks regarding object creation	231
10.2.1 Creation of table with XML columns	232
10.2.2 Alteration of implicitly created XML objects	232
10.2.3 Sizing table spaces	232
10.2.4 Compression	234
10.2.5 Registration of schemas	235
10.2.6 Creation of XML indexes	236
10.2.7 Grants and authorizations required	236
10.3 Housekeeping	237
10.4 Backup and recovery	237
10.5 Diagnostics	238
10.5.1 Identification of XML related objects	238
10.5.2 Investigating XML specific errors	238
10.5.3 Correcting XML data	240
Chapter 11. Performance considerations	243
11.1 Choice of relational or XML storage	244
11.1.1 XML only storage	244
11.1.2 Hybrid storage	245
11.1.3 Natural fit for XML storage	246
11.2 XML Schema validation	248
11.3 Managing access path selection with XML	249
11.3.1 Differences between XML and relational indexes	249
11.3.2 XML index design	250
11.4 Encourage use of native SQL DB2 routines	256
11.5 External language programming	257
11.6 DBA considerations	257
11.7 SQL/XML coding techniques	260
11.7.1 XMLTABLE to minimize database calls	260
11.7.2 XMLEXISTS for index access	261
11.7.3 Simple XPath expressions	261
Appendix A. Application scenario documents	263
A.1 Schema	264
A.2 XML message	264
Appendix B. Additional material	273
Locating the Web material	273
Using the Web material	273
System requirements for downloading the Web material	273
Downloading and extracting the Web material	274
Glossary	279
Related publications	293
IBM Redbooks	293
Other publications	293
Online resources	293

How to get Redbooks	294
Help from IBM	294
Index	295

Figures

1-1 XML: The foundation for Web services	3
1-2 DOM tree	12
1-3 DOM tree after XSL transformation	13
1-4 Hello world	19
2-1 SQL/XML query with XMLEXISTS predicate	27
2-2 XMLTABLE function example	29
2-3 z/OS XML system services and zAAP processing flow	42
2-4 Splitting an XPath expression over multiple lines in 3270 SPUFI session	43
3-1 Application scenario	47
3-2 Four application code samples	49
4-1 XML objects for segmented base table space	53
4-2 XML objects for partition-by-growth base table space	53
4-3 XML objects for classic partitioned base table space	55
4-4 XML objects for range-partitioned base table space	55
4-5 XML basic storage scheme	57
4-6 XML multi-versioning scheme	60
4-7 Multi-versioning for XML data	61
4-8 XML Locking scheme with multi-versioning	62
4-9 -DISPLAY DATABASE command output	67
5-1 XML schemas	76
5-2 XML Schemas in XML Schema Repository	79
5-3 Schema determination	79
6-1 SQL Query - SELECT * FROM BK_TO_CSTMTR_STMT	100
6-2 Optim Development Studio XML document viewer - Design view	101
6-3 Optim Development Studio XML document viewer - Source view	102
6-4 Tabular result set of bank statement entries	103
6-5 Relational result set spanning data elements from multiple XML documents	104
6-6 Typical "Ntry" node within a Bk_To_Cstmr_Stmt document	105
6-7 SQL Results using UDFs on XML documents	114
6-8 Scenario to receive XML CDC messages into DB2 pureXML via MQ	121
6-9 Initial CUST_HISTORY table contents for 'CUST1'	122
6-10 Updated CUST_HISTORY table contents for 'CUST1'	125
6-11 DB2 PureXML as historical repository	127
7-1 Exchange data as textual or binary XML format	136
7-2 Bank To Customer Statement example	142
7-3 HTML output after XSLT	150
7-4 Put a message into a Queue	151
8-1 Data conversion in a three-layer structure using CLOBs or BLOBs	159
8-2 BankToCustomerStatement message as shown in a browser	162
8-3 Subset of the BankToCustomerStatement schema	163
8-4 Message recipient of a BankToCustomerStatement	167
8-5 Schema definition for the GrpHdr element	170
8-6 MsgRcpt element with namespace declaration	170
8-7 Schema definition of GrpHdr element	174
8-8 Revised schema definition for GrpHdr with multiple MsgRcpt elements	174
8-9 MsgRcpt element created with XML GENERATE	177
8-10 MsgRcpt element created with XML GENERATE WITH ATTRIBUTES	178
9-1 CHECK DATA syntax - new keywords	183

9-2	CHECK DATA - SHRLEVEL REFERENCE considerations	185
9-3	CHECK DATA - SHRLEVEL CHANGE considerations (1 of 2)	185
9-4	CHECK DATA - SHRLEVEL CHANGE considerations (2 of 2)	186
9-5	Make a partial update to the XML document	190
9-6	Copy enable the DOCID and NODEID indexes	204
9-7	Status of database DSN00242	207
9-8	Content of the XML document	208
9-9	Status of database DSN00242 (after partial recovery)	210
9-10	Status of database DSN00242 (after partial recovery and rebuild of indexes)	210
9-11	Content of the XML document after partial recovery	211
11-1	XML-only database design	244
11-2	Hybrid storage model	245
11-3	Physical access path using an XML Index.	250
11-4	Catalog query to sysibm.sysindexes.	252

Tables

1-1 XPath	14
2-1 IBM tools for DB2 administration and development with DB2 pureXML	44
4-1 Data in table T1	58
6-1 MQ scalar functions provided by DB2 10	94
6-2 IBM change data capture tools that publish XML messages	119
7-1 XML data type support in JCC3 and JCC4	130
7-2 JDBC 3.0 Getter methods of ResultSet	131
7-3 DB2Xml Getter Methods	131
7-4 Methods to retrieve XML data from SQLXML object	133
7-5 Method to set XML value to SQLXML object	133
8-1 Insert an XML document with different host variable types	158
8-2 Options for file reference variables	160
9-1 CHECK DATA invocation	184
9-2 Action to be taken based on CHECK INDEX output	188
9-3 Example LISTDEF statements	193
10-1 Properties can be altered for XML objects	232
10-2 Table space types for base and XML tables	233
10-3 DSSIZE of the XML table space	233
10-4 Restricted states related to XML	239
10-5 Corrective action after running CHECK INDEX	240

Examples

1-1	An XML document	5
1-2	DTD	6
1-3	XML document	6
1-4	Introducing need for namespaces	7
1-5	The need for namespaces	8
1-6	Namespaces	8
1-7	Namespaces without prefixes	9
1-8	A book description	10
1-9	XML Schema	10
1-10	XPath	14
1-11	Sample XQuery	15
1-12	Strict DTD	16
1-13	Transitional DTD	16
1-14	Frameset DTD	16
1-15	Simple tags	17
1-16	hello.xml	18
1-17	hellohtm.xsl	18
1-18	Output from XSLT processor	18
2-1	Defining and populating an XML column	24
2-2	XML Publishing Functions of SQL/XML	25
2-3	Example table with XML column	25
2-4	Simple XMLEXISTS example	26
2-5	XMLEXISTS example with namespace declaration	28
2-6	Simple XMLTABLE example	28
2-7	Simple XMLQUERY example	30
2-8	XMLCAST example casting a numeric data element to varchar or integer	30
2-9	XMLPARSE function and whitespace handling	31
2-10	XMLSERIALIZE example to convert an XML document to a UTF-8 CLOB	31
2-11	Hybrid data access example	32
2-12	Initial contents of XMLADDRESS table	33
2-13	XMLMODIFY to replace a node	33
2-14	XMLMODIFY to delete a node	33
2-15	XMLMODIFY to insert a node	34
2-16	XML Index creation examples	35
2-17	Simple XML schema example	36
2-18	XML document that conforms to previous XML schema	37
2-19	XML schema registration	37
2-20	SYSXSR.DSN_XMLVALIDATE example	37
2-21	XML type modifier example	37
2-22	DSNTIJRV installation verification job output	39
2-23	z/OS console display WLM APPLENV status	40
2-24	Check XSR tables exist	40
2-25	Check XSR routines exist	40
2-26	Creation of Java stored procedures	41
4-1	BK_TO_CSTMR_STMT table with XML column	52
4-2	Catalog Queries (1 of 3)	63
4-3	Catalog Queries (2 of 3)	65
4-4	Catalog Queries (3 of 3)	66

4-5	Using the SQL INSERT statement to insert XML document to an XML column.	67
4-6	XML index on DiTm elements.	70
5-1	Specify an XML type modifier for an XML column at create time	76
5-2	Table definition without XML type modifier	76
5-3	Specify XML type modifier for XML column at alter time	76
5-4	Add an XML schema to the XML type modifier.	77
5-5	Reset XML type modifier for XML column at alter time.	77
5-6	Identify an XML schema by target namespace and schema location.	77
5-7	Identify an XML schema by target namespace.	78
5-8	No namespace	78
5-9	Specifying global element name	78
5-10	Schema selection for validation from an XML type modifier - Example 1	80
5-11	Schema selection for validation from an XML type modifier - Example 2.	81
5-12	Schema selection for validation from an XML type modifier - Example 3.	81
5-13	Schema selection for validation from an XML type modifier - Example 4.	81
5-14	Schema selection for validation for DSN_XMLVALIDATE - Example 1	82
5-15	Schema selection for validation for DSN_XMLVALIDATE - Example 2	83
5-16	Schema selection for validation for DSN_XMLVALIDATE - Example 3	83
5-17	Schema selection for validation for DSN_XMLVALIDATE - Example 4	83
5-18	Search for documents not validated	84
5-19	Retrieve target namespaces and XML schema names used for validation	85
6-1	Tables used for following examples.	88
6-2	Registering the ISO20022 Bnk_To_Cst_Stmt XML schema.	89
6-3	Simple stored procedure for registering the ISO20022 Bnk_To_Cst_Stmt XML schema. 89	
6-4	Stored procedure with error handling logic	91
6-5	Populating the MQSERVICE-TABLE	94
6-6	Sample MQREAD and MQSEND function calls	95
6-7	Stored procedure to read XML message from MQ	96
6-8	Command to configure MQ listener	97
6-9	Command to show MQ listener configuration.	97
6-10	Contents of SYSMQL.LISTENERS.	98
6-11	Stored procedure modified for MQ listener integration	98
6-12	Commands to operate MQ listener	99
6-13	SQL/XML Query to yield a “traditional” style bank statement	102
6-14	SQL/XML query spanning multiple XML documents	103
6-15	Candidate XML index definitions	106
6-16	Calculating the sum of the entries in a BankToCustomerStatement.	107
6-17	Using time and date functions in an XPath expression.	108
6-18	Avoiding empty sequences in result by using XMLEXISTS	109
6-19	Combining XMLQUERY with publishing functions	109
6-20	Using XMLAGG to consolidate all entries into one document	110
6-21	Extracting one entry per row using XMLTABLE	111
6-22	Grouping entries obtained from XMLTABLE according to currency	112
6-23	Creating three user defined functions	113
6-24	Usage of UDFs	113
6-25	Modified stored procedure using UDFs instead of XQuery expressions.	114
6-26	UDF for XML sub-document update	115
6-27	Contents of relational address table	116
6-28	XML to Relational Join example using XMLEXISTS.	116
6-29	XML to relational join example using XMLTABLE	117
6-30	Script to convert the relational address table to XML	117
6-31	XML to XML join using XMLEXISTS.	118

6-32	Sample XML CDC message format for DB2 and Classic Data Event Publishers. . .	119
6-33	Sample XML CDC message format for InfoSphere CDC	120
6-34	Stored Procedure to receive and apply CDC message	122
6-35	XMLTABLE function for Event Publisher XML schemas.	125
7-1	Creating a GroupHeader of the BankToCustomerStatement message	133
7-2	Constructing XML as a DOM tree	134
7-3	Setting the xmlFormat.	136
7-4	Register XML schema in Java application	138
7-5	DDL for table BK_TO_CSTMR_STMT	139
7-6	Parsing XML value and inserting into DB2	139
7-7	Modifying an XML document	142
7-8	Retrieving the entire or partial XML document	144
7-9	DDL for STMT table	145
7-10	Creating a SQLstored procedure	145
7-11	Handling the SQL stored procedure	147
7-12	Expecting output after transform.	148
7-13	XSLT example to transform from XML to XML	148
7-14	Java application to transform XML document.	149
7-15	XSLT example to transform from XML to HTML.	149
7-16	Java example to put a message into a queue	152
8-1	XML host variables in COBOL	156
8-2	XML host variables after transformation by the DB2 pre-compiler	156
8-3	XML declaration with encoding declaration as an attribute.	157
8-4	Explicit declaration of variable CCSID	157
8-5	XML file reference filterable in COBOL.	160
8-6	XML file reference variables after transformation by the DB2 pre-compiler	160
8-7	Initialization of a file reference variable.	161
8-8	CLP commands for registration of XML schema.	164
8-9	DDL for the table in the BankStmt application	164
8-10	Insert a BankToCustomerStatement.	164
8-11	Extracting key fields using XMLTABLE.	165
8-12	JCL for running COBOL insert program	165
8-13	Validation error on insert.	166
8-14	Determining whether an XML document has been validated	166
8-15	COBOL program for updating a BkToCstmrStmt with a new MsgRcpt.	168
8-16	SQL error when updating XML document with MsgRcpt element	169
8-17	Retrieval of an XML document to a file	171
8-18	Retrieval of data in relational format from an XML document.	172
8-19	Candidate index pattern for the BankStmt application	173
8-20	XML index on DtTm elements.	173
8-21	Access path using the index IXMLNTRY	173
8-22	Adding a new schema to an XML type modifier	175
8-23	Insert a MsgRcpt element after the CreDtTm element	175
8-24	COBOL program for generation of the MsgRcpt element.	176
8-25	COBOL program for shredding a MsgRcpt element into variables.	178
8-26	Converting a schema to OSR format	179
8-27	XMLPARSE with schema validation	179
8-28	DD statement for supplying a schema to the COBOL program	180
9-1	CHECK DATA example	184
9-2	CHECK INDEX utility JCL and output.	186
9-3	COPY utility JCL for taking full image copy and output	188
9-4	COPY utility JCL for taking incremental image copy and output.	190
9-5	COPYTOCOPY utility JCL and output	191

9-6	JCL for LISTDEF utility and output (1 of 3)	193
9-7	JCL for LISTDEF utility and output (2 of 3)	194
9-8	JCL for LISTDEF utility and output (3 of 3)	195
9-9	LOAD utility JCL and output	196
9-10	LOAD utility JCL (using file reference variable) and output	198
9-11	LOAD utility JCL and output (input to load is in binary format)	199
9-12	LOAD utility JCL and output (input to load is in spanned record format)	201
9-13	MERGECOPY utility JCL and output	202
9-14	QUIESCE utility JCL and output (1 of 2)	203
9-15	QUIESCE utility JCL and output (2 of 2)	204
9-16	REBUILD INDEX utility JCL and output	205
9-17	RECOVER TABLESPACE utility JCL and output	208
9-18	RECOVER TABLESPACE utility JCL (with modified control statement) and output	208
9-19	REORG TABLESPACE utility JCL and output	213
9-20	REPORT utility JCL (with TABLESPACESET option) and output	215
9-21	REPORT utility JCL (with RECOVERY option for base table space) and output	216
9-22	REPORT utility JCL (with TRECOVERY option for XML table space) and output	217
9-23	RUNSTATS utility JCL and output	219
9-24	UNLOAD utility JCL and output	221
9-25	UNLOAD utility JCL (using file reference variable) and output	222
9-26	UNLOAD utility JCL (to unload XML data in binary) and output	223
9-27	UNLOAD utility JCL (to unload XML data in spanned record format) and output	224
9-28	DSNTIAUL with SQL parameter	225
9-29	DSNTIAUL with LOBFILE parameter	227
10-1	Creating a range-partitioned table	233
10-2	Creating an XML index with compression	235
10-3	Create an XML index	236
10-4	Display database command shows XML table space in AUXW	238
10-5	REPAIR LOCATE control statements for diagnosing XML inconsistencies	240
10-6	Using REPAIR utility to clear ACHKP status on table space	241
11-1	A lean XML index	251
11-2	A heavy XML index	251
11-3	A “silly” XML index	251
11-4	Explain for XMLEXISTS	253
11-5	Explain for XMLTABLE with XMLEXISTS	254
11-6	Explain for XMLTABLE with an XML predicate	255
11-7	Explain for XMLQUERY with XMLEXISTS	255
11-8	sysibm.syscolumns contents for auxiliary XML table space	257
11-9	REORG of a table space with an XML table space	259
11-10	Multiple XMLQUERY calls replaced with a single XMLTABLE call	260
11-11	single select statement combining two xmlquery expressions	260
A-1	XML message received	264
A-2	XML message parts processed	269

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	IMS™	RACF®
CICS®	Informix®	Rational®
DataPower®	InfoSphere™	Redbooks®
DB2 Connect™	iSeries®	Redpaper™
DB2®	MVS™	Redbooks (logo)  ®
Domino®	Optim™	S/390®
DRDA®	OS/390®	System z®
ESCON®	OS/400®	VisualAge®
FICON®	PR/SM™	WebSphere®
IBM®	pureXML®	z/OS®

The following terms are trademarks of other companies:

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows NT, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

The DB2® pureXML® feature offers sophisticated capabilities to store, process and manage XML data in its native hierarchical format. By integrating XML data intact into a relational database structure, users can take full advantage of DB2's relational data management features.

In this IBM® Redbooks® publication we document the steps for the implementation of a simple but meaningful XML application scenario. We have chosen to provide samples in COBOL and Java™ language. The purpose being to provide an easy path to follow to integrate the XML data type for the traditional DB2 user.

We also add considerations for the data administrator and suggest best practices for ease of use and better performance.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Paolo Bruni is a DB2 Information Management Project Leader at the International Technical Support Organization based in the Silicon Valley Lab. He has authored several IBM Redbooks publications about DB2 for z/OS® and related tools and has conducted workshops and seminars worldwide. During Paolo's many years with IBM, in development and in the field, his work has been mostly related to database systems.

Neale Armstrong is a Consultant IT Specialist in IBM United Kingdom, responsible for technical support for System z® Information Management products. He has 24 years of experience in DB2 solutions on z/OS and distributed platforms. He holds a degree in Physics from the University of Bristol. His areas of expertise include database federation, replication and event publishing, for DB2, IMS™ and VSAM data sources, which could more generally be referred to a “database plumbing”. He has co-authored three previous IBM Redbooks publications.

Ravi Kumar is a Senior Instructor and Specialist for DB2 with IBM Software Group, Australia. He has about 25 years of experience in DB2. He was on assignment at the International Technical Support Organization, San Jose Center, as a Data Management Specialist from 1994 to 1997. He has co-authored many IBM Redbooks publications including *DB2 UDB for z/OS Version 8 Everything You Ever Wanted to Know, ... and More*, SG24-6079, *DB2 9 for z/OS Technical Overview*, SG24-7330, and *DB2 10 for z/OS Technical Overview*, SG24-7892. He is currently on virtual assignment as a Course Developer in the Education Planning and Development team, Information Management, IBM Software Group, USA.

Kirsten Ann Larsen is a senior IT specialist and technical lead with IT Delivery in IBM Nordics. She has 14 years of experience with DB2 for z/OS and has co-authored the IBM Redbooks publication *Securing DB2 and Implementing MLS on z/OS*, SG24-6480. She holds a master's degree in Computer Science from Aarhus University. She has worked with XML since pureXML support was included with the release of DB2 9 in 2007 and has co-authored a number of articles on XML.

Tink Tysor is president of Bayard Lee Tysor, Inc. in the United States of America. He has 40 years of experience in programming, the last 14 years specializing in DB2. He holds a degree in Economics from American University. His areas of expertise include SQL, DB2 Database Administrating, and Data Modeling. He has written and presented extensively on SQL for DB2 including the IBM Redbooks publication *DB2 for z/OS Tools for Database Administration and Change Management*, SG24-6480-00.

Hao Zhang is a software engineer in IBM China Software Development Lab. He has over 6 years of experience in DB2 QA field. He participated in testing several DB2 pureXML features in DB2 9 and DB2 10 for z/OS, and presented DB2 9 pureXML support in CDUG (Chinese DB2 Users' Group) in 2009. His areas of expertise include distributed area in JCC driver, temporal table and XML.

Thanks to the following people for their contributions to this project:

Rich Conway
Bob Haimowitz
Emma Jacobs
International Technical Support Organization

Li Chen
Mengchu Cai
Jason Cu
Thanh Dao
Eric Katayama
Andrew Lai
Susan Malaika
Gary Mazo
Roger Miller
Jinfeng Ni
Matthias Nicola
Bryan Patterson
Tom Ross
Guogen Zhang
IBM Silicon Valley Lab

Heidi Arnold
IBM Boeblingen

Michael Schwartzbach
Aarhus University

Rick Butler
BMO Toronto

Lee Ackerman
IBM Ottawa

Nagesh Subrahmanyam
IBM India

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your

area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>



Introduction

This chapter provides an introduction to XML technology, its importance in the IT business, and the contents of the book.

This chapter contains the following:

- ▶ Importance of XML data
- ▶ XML introduction
- ▶ What is in this book

1.1 Importance of XML data

XML technology has become pervasive in virtually all industries and sectors, owing to its versatility and neutrality for exchanging data among diverse devices, applications, and systems from different vendors. These qualities of XML along with its easy to understand self-describing nature, ability to handle structured, semi-structured and unstructured data, and support for Unicode have made XML a universal standard for data interchange.

1.1.1 Growth of XML

Nearly every company today comes across XML in some form. The amount of XML data that organizations have to deal with is growing at a rapid rate. In fact, the volume of XML data is growing faster than the traditional data that typically resides in relational databases. Factors fueling the growth of XML data include:

- ▶ XML-based industry and data standards
- ▶ Service-oriented architectures (SOA) and Web services
- ▶ Web 2.0 technologies such as XML feeds and syndication services

XML-based industry standards

Almost every industry has multiple standards based on XML and there are numerous cross-industry XML standards as well. A few examples of XML-based industry standards are listed here:

- ▶ ACORD - XML for the Insurance Industry:
<http://www.acord.org/>
- ▶ FPML - Financial Product:
<http://www.fpml.org/>
- ▶ HL7 - Health Care:
<http://www.hl7.org/>
- ▶ IFX - Interactive Financial Exchange:
<http://www.ifxforum.org/>
- ▶ IXRetail - Standard for Retail operation:
<http://www.nrf-arts.org/>
- ▶ XBRL - Business Reporting / Accounting:
<http://www.xbrl.org/>
- ▶ NewsML - News / Publication:
<http://www.newsml.org/>

These standards facilitate purposes such as the exchange of information between the various players within these industries and their value chain members, data definitions for ongoing operations, and document specifications. More and more companies are adopting such XML standards or are being compelled to adopt them in order to stay competitive, improve efficiencies, communicate with their trading partners or suppliers, or just to perform everyday tasks.

SOA and Web services

Service-oriented frameworks and deployments are growing in popularity, owing to their ability to integrate systems, permit reuse of resources, and respond quickly to changing market

conditions, allowing companies to save money and improve competitiveness. In service-oriented architectures, consumers and service providers exchange information using messages. These messages are invariably encapsulated as XML. As such, XML can provide the plumbing in SOA environments as illustrated in Figure 1-1. Therefore the drive towards information as a service and rapid adoption of SOA environments is also stimulating the growth of XML.

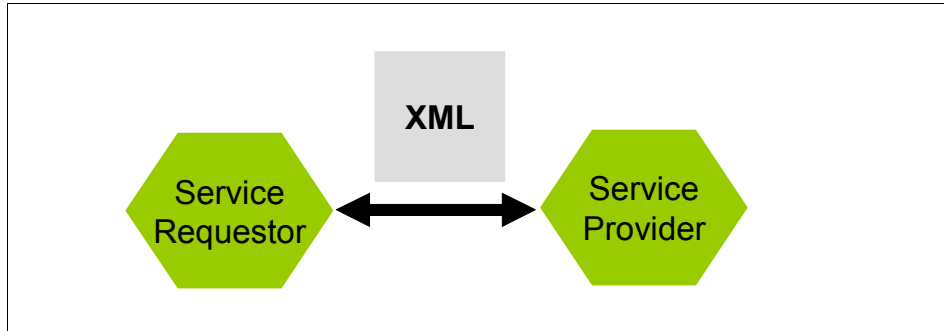


Figure 1-1 XML: The foundation for Web services

Web 2.0 technologies

Syndication is considered to be the heartbeat of Web 2.0, the next generation of the Internet. Atom and RSS feeds can be found abundantly on the Web, allowing the user to subscribe to them and be kept up-to-date about all kinds of Web content changes, such as news stories, articles, wikis, audio and video files.

Content for these feeds is rendered as XML files and can contain links, summaries, full articles, and even attached multimedia files such as podcasts. Syndication and Web feeds are transforming the Web as we know it. New business models are emerging around these technologies. As a consequence, XML data now exists not only in companies adopting XML industry standards, or enterprises implementing SOAs, but also on virtually every Web-connected desktop.

1.1.2 The value of XML data

As a result of XML industry standards becoming more prevalent, the drive towards SOA environments, and rapid adoption of syndication technologies, more and more XML data is being generated every day as Web feeds, purchase orders, transaction records, messages in SOA environments, financial trades, insurance applications, and other industry-specific and cross-industry data. That is, XML data and documents are becoming an important business asset containing valuable information (such as customer details, transaction data, order records, and operational documents).

The growth and pervasiveness of XML assets presents challenges and opportunities for companies. When XML data is harnessed, and the value of the information it contains is unlocked, it can translate into opportunities for organizations to streamline operations, derive insight, and become agile.

On the other hand, as XML data becomes more critical to the operations of an enterprise, it presents challenges in that XML data must be secured, maintained, searched, and shared. Depending on its use, XML data might also have to be updated, audited, and integrated with traditional data. All of these tasks must be done with the reliability, availability, and scalability afforded to traditional data assets.

That is, in order to unleash the potential of XML data, it requires storage and management services similar to what enterprise-class relational database management systems such as DB2 have been providing for relational data.

1.2 XML introduction

This brief introduction to XML is extracted from *XML on z/OS and OS/390: Introduction to a Service-Oriented Architecture*, SG24-6826.

The idea of universal data formats is not new. Programmers have been trying to find ways to exchange information between different computer programs for a long time. Standard Generalized Markup Language (SGML) was developed to achieve this. SGML can be used to *mark up data*, that is, to add metadata in a way that allows data to be *self-describing*. SGML is meta-language.

The markup process involves using *tags* to identify pieces of information in a document. Tags are names (strings of characters) surrounded by arrow brackets (< and >). Every piece of data that is encoded will have a start tag and an end tag, for example, <town> patiya</town>. The start and end tags make it easy for software to process the encoded information, as it clearly delineates where certain pieces of information start and where they end.

SGML does not prescribe any particular markup; instead, it defines how any markup language can be formally specified.

The most popular SGML application is HTML (Hypertext Markup Language), the markup language that rules the Web. The HTML specification is owned by W3C. However, different browser vendors introduced a number of incompatible tags to HTML, which are outside the scope of the original HTML specifications. These tags create problems for developers when they author Web pages because they must consider what browser will display the pages. And, although HTML has been very successful for displaying information on browsers, it was not found to be useful in describing the data that it represents, meaning it did not have the metadata capability that is essential for a self-describing data document.

Furthermore, SGML is quite inefficient and cumbersome when it is used to encode complex data structure. Hence, there arose a need to develop a more lightweight markup language, so W3C developed the specification for XML (eXtensible Markup Language). XML is similar to SGML in that it preserves the notion of *general markup*. There are very few optional features, and most SGML features that were deemed difficult to implement have been dropped.

In this section we look at the following topics:

- ▶ XML definitions
- ▶ Document validity and well-formedness
- ▶ XML Schema
- ▶ Extensible Stylesheet Language
- ▶ XPath
- ▶ XQuery
- ▶ XHTML
- ▶ XSL, XSLT, Xpath, and XHTML examples

1.2.1 XML definitions

XML is a system-independent standard for the representation of data. XML is not just some new version of HTML; it is different from HTML. Like HTML, XML has tags, and in these tags

it encloses data. Where an HTML tag says something like “display this data in bold font” (...), an XML tag acts like a field name in your program. It puts a label on a piece of data that identifies it (for example: <message>...</message>). This is the first of a number of differences between the languages.

In XML you can create the tags you want, with only a small number of restrictions, and these tags are used by a program (parser) to process the data enclosed between them.

Example 1-1 shows a simple XML document.

Example 1-1 An XML document

```
<?xml version="1.0"?>
<!DOCTYPE JavaXML:EmployeeList SYSTEM "DTD\JavaXML.dtd">
<JavaXML:employeeList xmlns:JavaXML="http://www.ibm.com">
  <JavaXML:Employee action="add">
    <JavaXML:firstName>David</JavaXML:firstName>
    <JavaXML:secondName>Sanchez Carmona</JavaXML:secondName>
    <JavaXML:age>20</JavaXML:age>
  </JavaXML:Employee>
  <JavaXML:Employee action="delete">
    <JavaXML:firstName>Jose Luis</JavaXML:firstName>
    <JavaXML:secondName>Fernandez Lastra</JavaXML:secondName>
  </JavaXML:Employee>
</JavaXML:employeeList>
```

Clients with their Web browser could fill out a form, entering the names of the employees they want to add or delete. The data could then be sent to a Web application that could process the XML document and extract the data, generating the necessary updates, for example, on a DB2 table.

As this example illustrates, the rules are very few: each tag must have an enclosing tag, and not much more. The tags are invented tags, which means that they are free-form.

Text is system-independent, and since XML is very flexible and is based only on text, it is used as the main way to transport data between different environments.

Often, XML documents are automatically generated by tools, and in many situations we need these XML documents to follow rules we create. We use other documents, containing XML data definitions in which we specify our restrictions, to accomplish this.

Document Type Definition (DTD) is a set of markup declarations that define a document type for SGML-family markup languages (SGML, XML, HTML). DTD is described in “Document Type Definition” on page 7.

*XML Schema*¹ is another rules language that aims to provide more complex semantic rules. It also introduces new semantic capabilities, such as support for namespaces and type-checking within an XML document. XML Schema is described in 1.2.3, “XML Schema” on page 9.

1.2.2 Document validity and well-formedness

XML documents can be *well-formed*, or they can be *well-formed and valid*. These are two very important rules that do not exist for HTML documents. These iron-clad rules contrast

¹ The W3C-recommended schema language for XML is XML Schema, see <http://www.w3.org/XML/Schema>.

with the more free-style nature of a lot of the concepts in XML. The rules can be defined briefly as follows:

- ▶ A well-formed document satisfies a list of syntax rules provided in the specification for XML documents.
- ▶ A valid document contains a reference to a Document Type Definition (DTD) or XML Schema Definition (XSD), and its elements and attributes follow the grammatical rules that the DTD or XSD specifies.

Schema languages typically constrain the set of elements that may be used in a document, which attributes may be applied to them, the order in which they may appear, and the allowable parent/child relationships

XSD schema, often referred to as (XML Schema, is a newer schema language, successor to DTD language. XSD documents are far more powerful than DTD's in describing XML languages. They use a rich datatyping system and allow for more detailed constraints on an XML document's logical structure. XSDs also use an XML-based format, which makes it possible to use ordinary XML tools to help process them. This has become the more popular approach to working with XSD.

With few exceptions, every DTD can be converted to an equivalent XML Schema.

Difference between well-formedness and validity

All of the constraints are defined in the XML 1.0 recommendation. For more information refer to the Web site:

<http://www.w3.org/XML>

Determining whether a particular document is in compliance with these rules is a two step process. Well-formedness insures that XML parsers is able to read the document, validity (which implies well-formedness) determines whether an XML document adheres to a DTD or XML Schema. An XML application checks for and rejects documents that are not well-formed before checking whether they comply with validity constraints (VCs).

A document might be well-formed but still not be valid. The following examples illustrate the difference between well-formedness and validity:

- ▶ Documents that adhere to rules described in the associated DTD or XSD are valid.
- ▶ Documents that carry out the syntactical rules for XML documents are well-formed. These rules have to do with attribute names, which should be unique within an element, and attribute values, which must not contain the character <, and so on.

Example 1-2 shows a DTD.

Example 1-2 DTD

```
<!ELEMENT BANCO (TARJETA+)>
<!ELEMENT TARJETA (Nom, Cod_Cuenta)>
<!ELEMENT Nom (#PCDATA)>
<!ELEMENT Cod_Cuenta ((#PCDATA)>
```

Example 1-3 shows a sample XML document.

Example 1-3 XML document

```
<BANCO>
  <TARJETA>
    <NOM>Silvia</NOM>
```

```

    <Cod_Cuenta>2562789452</Cod_Cuenta>
  </TARJETA>
</BANCO>

```

The document shown in Example 1-3 is well-formed, but it is not valid according to the sample DTD shown in Example 1-2 because the `<NOM>` tag is not defined in the associated DTD (tags are case sensitive).

Document Type Definition

A Document Type Definition, or DTD, specifies the kinds of tags that can be included in your XML document, the valid arrangements of those tags, and the structure of the XML document. The DTD defines the type of elements, attributes, and entities allowed in the documents, and may also specify some limitations to their arrangement. You can use the DTD to make sure you don't create an invalid XML structure since the DTD defines how elements relate to one another within the document's tree structure. You can also use it to define which attributes can be used to define an element and which ones are not allowed.

The DTD can be either stored in a separate file or embedded within the same XML file. If it is stored in a separate file it may be shared with other documents.

An XML document is not required to have a DTD. DTDs provide parsers with clear instructions on what to check for when they are determining the validity of an XML document. DTDs or other mechanisms, like XML schemas, contribute to the goal of ensuring that the application can easily determine whether the XML document adheres to a given set of rules, beyond the well-formedness rules defined in the XML standard.

DTDs do have limitations, for example:

- ▶ A DTD makes it possible to validate the structure of relatively simple XML documents, but that's as far as it goes. A DTD can't restrict the content of elements, and it cannot specify complex relationships.
- ▶ In a DTD, you only get to specify the structure of the `<heading>` element one time. There is no context-sensitivity.
- ▶ A DTD specification is not hierarchical. For a mailing address that contains several "parsed character data" (PCDATA) elements, for instance, the DTD might look something like that shown in Example 1-4. As you can see, the specifications are linear. That fact forces you to come up with new names for similar elements in different settings.

Example 1-4 Introducing need for namespaces

```

<!ELEMENT mailAddress (name, address, zipcode)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT address (#PCDATA)>
<!ELEMENT zipcode (#PCDATA)>

```

- ▶ Another problem with the non-hierarchical nature of DTD specifications is that it is not clear what comments are meant to explain.
- ▶ Finally, a DTD uses syntax which is substantially different from XML, so it cannot be processed with a standard XML parser. That means you cannot read a DTD into a DOM², for example, modify it, and then write it back out again.

² XML Document Object Models (DOM) is a programming interface for HTML and XML documents which defines the way a document can be accessed and manipulated.

Namespaces

Before talking about XML Schema, we must first clarify the concept of *Namespaces*. Namespaces are used when there is a need to have different elements with different attributes but with the same name. Depending on the context, a tag is related to an element or to another one. Example 1-5 illustrates this situation.

Example 1-5 The need for namespaces

```
<widget type="gadget">
  <head size="medium"/>
    <info>
      <head>
        <title>Description of gadget</title>
      </head>
      <body>
        <h1>Gadget</h1>
      </body>
    </info>
  </widget>
```

It is obvious that there is a problem with the meaning of `<head>`. It depends on the context. This situation complicates things for processors and might even cause ambiguities. We need some mechanism to distinguish between the two, and apply the correct semantic description to the correct tag. The root of the problem is one common name space.

There is a simple solution to this problem: namespaces. Namespaces are a simple and straightforward way to distinguish names used in XML documents. If you can specify the related DTD when an element is being validated, the problem is solved.

As you can see in Example 1-6, the `<title>` tag is used twice, but in a different context: once within the `<author>` element and once within the `<book>` element. Note the use of the `xmlns` keyword in the namespace declaration, one for *authr*, one for *bk*. Interestingly, the XML recommendation does not specify whether a namespace declaration should point to a valid Uniform Resource Identifier (URI), only that it should be unique and persistent.

Example 1-6 Namespaces

```
<?xml version="1.0" ?>
<library-entry xmlns:authr="http://sc58ts.itso.ibm.com/Jose/2002/author.dtd"
xmlns:bk="books.dtd">
  <bk:book>
    <bk:title>XML Sample</bk:title>
    <bk:pages>210</bk:pages>
    <bk:isbn>1-868640-34-2</bk:isbn>
    <authr:author>
      <authr:firstname>JuanJose</authr:firstname>
      <authr:lastname>Hernandez</authr:lastname>
      <authr:title>Mr</authr:title>
    </authr:author>
  </bk:book>
</library-entry>
```

In Example 1-6, in order to illustrate the relationship of each element to a given namespace, we specify the relevant namespace prefix before each element. Prefixes are bound to namespace URIs by attaching the *xmlns:prefix* attribute to the prefixed element or one of its ancestors. Bindings have scope within the element where they are declared.

Once a prefix is applied to an element name, it applies to all descendants of that element unless it is overridden by another prefix. The extent to which a namespace prefix applies to elements in a document is defined as the namespace scope.

Example 1-7 is equivalent to Example 1-6, but only the necessary namespace prefixes have been used.

Example 1-7 Namespaces without prefixes

```
<?xml version="1.0" ?>
<library-entry xmlns:authr="http://sc58ts.itso.ibm.com/Jose/2002/author.dtd"
xmlns:bk="books.dtd">
  <bk:book>
    <title>XML Sample</title>
    <pages>210</pages>
    <isbn>1-868640-34-2</isbn>
    <authr:author>
      <firstname>JuanJose</firstname>
      <lastname>Hernandez</lastname>
      <title>Mr</title>
    </authr:author>
  </bk:book>
</library-entry>
```

Information on namespaces can be found at the following Web site:

<http://www.w3.org/TR/REC-xml-names>

1.2.3 XML Schema

The W3C XML Schema Definition (XSD) language is an XML language for describing and constraining the content of XML documents. A *Schema* is similar to a DTD in that it defines which elements an XML document can contain, how they are organized, and which attributes and attribute types elements can be assigned. Therefore it is a method to check the validity of well-formed XML documents. For more information, see <http://www.w3.org/XML/Schema>.

DTD and XML Schema

In “Document Type Definition” on page 7 we introduced DTD and identified some of its limitations. In addition to those limitations, we can add the following:

- ▶ There are no constraints on character data. If character data is allowed, any character data is allowed.
- ▶ The attribute value models are too simple.
- ▶ There is no support for namespaces.
- ▶ There is no support for schema evolution, extension, or inheritance of declarations.
- ▶ It is difficult to write, maintain, and read large DTDs, and to define families of related schemas.
- ▶ It is only possible to set defaults for attributes, not for elements.

Therefore, there is a need for a way to specify more complex semantic rules and provide all those things that DTDs cannot do, like type-checking within an XML document. XML Schema provides such functionality; it also introduces semantic capabilities, such as support for namespaces and type-checking.

The main advantages of XML Schemas over DTDs are:

- ▶ Schemas use XML syntax.

- ▶ It is possible to specify data types.
- ▶ Schemas are extensible.

XML Schema example

It is difficult to give a general outline of the elements of a schema due to the number of elements that can be used according to the W3C XML Schema Definition Language. The purpose of this language is to provide an inventory of XML markup constructs with which to write schemas. Example 1-8 is a simple document which describes the information about a book.

Example 1-8 A book description

```
<?xml version="1.0" encoding="UTF-8"?>
<book isbn="0836217462">
  <title>
    Don Quijote de la Mancha
  </title>
  <author>De Cervantes Saavedra, Miguel</author>
  <character>
    <name>Sancho Panza</name>
    <friend-of>El Quijote</friend-of>
    <since>1547-10-04</since>
    <qualification> escudero </qualification>
  </character>
  <character>
    <name>ElbaBeuno</name>
    <since>1547-08-22</since>
    <qualification>Amor Platonico de Don Quijote</qualification>
  </character>
</book>
```

Since the XML Schema is a language, there are several choices to build a possible schema that covers the XML document. Example 1-9 is a possible and very simple design.

Example 1-9 XML Schema

```
<?xml version="1.0" encoding="utf-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:element name="book">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="title" type="xs:string"/>
        <xs:element name="author" type="xs:string"/>
        <xs:element name="character" minOccurs="0" maxOccurs="unbounded">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="name" type="xs:string"/>
              <xs:element name="friend-of" type="xs:string" minOccurs="0"
                maxOccurs="unbounded"/>
              <xs:element name="since" type="xs:date"/>
              <xs:element name="qualification" type="xs:string"/>
            </xs:sequence>
          </xs:complexType>
        </xs:element>
      </xs:sequence>
      <xs:attribute name="isbn" type="xs:string"/>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

```
</xs:complexType>
</xs:element>
</xs:schema>
```

It is clear that Example 1-9 is an XML document since it begins with the XML document declaration. The schema element opens our schema holding the definition of the target namespace. Then we define an element named `book`. This is the root element in the XML document. We decided it is a complex type since it has attributes and non-text children. With sequence we begin to declare the children elements of the root element `book`. W3C XML Schema lets us define the type of data, as well as the number of possible occurrences of an element. For more information on possible values for these types, refer to the specification documents from W3C.

XML Schema options

XML Schema Language offers possibilities and alternatives beyond what is shown in Example 1-9. We could develop another schema based on a flat catalog of all the elements available in the instance document and, for each of them, lists of child elements and attributes. Thus we would have two choices: defining elements and attributes as they are needed, or creating them first and referencing them. The first option has a real disadvantage: the schema could become very difficult to read and maintain when documents are complex.

W3C XML Schema allows us to define data types and use these types to define our attributes and elements. It also allows the definition of groups of elements and attributes. In addition, there are several ways to arrange relationships between elements.

Documentation for XML Schemas can be defined by the `xs:documentation` element, and processing instructions for applications can be include with the `xs:appinfo` element.

As of August 2009, XSD 1.1 is a Candidate Recommendation with significant new features as defined at:

<http://www.w3.org/TR/xmlschema11-1/>
<http://www.w3.org/TR/xmlschema11-2/>

1.2.4 Extensible Stylesheet Language

Up to this point we have been concerned with XML—its syntax, how XML is used to mark up information according to our own vocabularies, how a program can check the validity of an XML document, and so forth. In other words, we have described how XML can ensure that an application running on any particular platform receives valid data. This is how we ensure that a program is able to process this data.

However, since XML only describes document syntax, the program does not know how to format this data without specific instructions about style.

The solution is XSL transformations. The Extensible Stylesheet Language (XSL) specification describes powerful tools to accomplish the required transformation of XML data. XSL consists of:

- ▶ The XSL Transformations (XSLT) language for transformation
- ▶ Formatting Objects (FO), a vocabulary for describing the layout of documents
- ▶ XSLT uses the XML Path Language (XPath), a separate specification that describes a means of addressing XML documents and defining simple queries.

XSLT offers a powerful means of transforming XML documents into other forms, producing XML, HTML, and other formats. It is capable of sorting, selecting, numbering, and has many other features for transforming XML. It operates by reading a style sheet, which consists of one or more templates, then matching the templates as it visits the nodes of the XML document. The templates can be based on names and patterns.

XSLT is increasingly being used to transform XML data into another form, sometimes different XML (for example, filtering out certain data, SQL statements, plain text, and so on), or any other format. Thus, any XML document may be shown in different formats, such as HTML, PDF, RTF, VRML, Postscript, and so forth.

The question is how to access and display the information contained in an XML file. After all, data is useless unless you can use it. This is where XSLT comes into the picture.

A comparison can be made between the relationship of CSS³ and HTML and the relationship of XSLT and XML. Indeed, XSLT is usually referred to as the *stylesheet language* of XML; however XML and XSLT are far more sophisticated technologies than HTML and CSS.

XSLT is a high-level declarative language. It is also a transforming and formatting language. It behaves in the following way:

- ▶ The pertinent data is extracted from an XML source document and transformed into a new data structure that reflects the desired output. The XSLT markup is commonly called a *stylesheet*. A parser is used to convert the XML document into a tree structure composed of various types of nodes. The transformation is accomplished with XSLT by using pattern matching and templates. Patterns are matched against the source tree structure, and templates are used to create a result tree.
- ▶ Next, the new data structure is formatted, for example in HTML or as text, and finally the data is ready for display.

Figure 1-2 shows the source tree from the XML document shown in Example 1-10.




Figure 1-2 DOM tree

The result tree after an XSL transformation could be an XHTML document, as shown in Figure 1-3.

³ Cascading Style Sheets (CSS) is a simple mechanism for adding style (e.g., fonts, colors, spacing) to Web documents. See <http://www.w3.org/Style/CSS/>

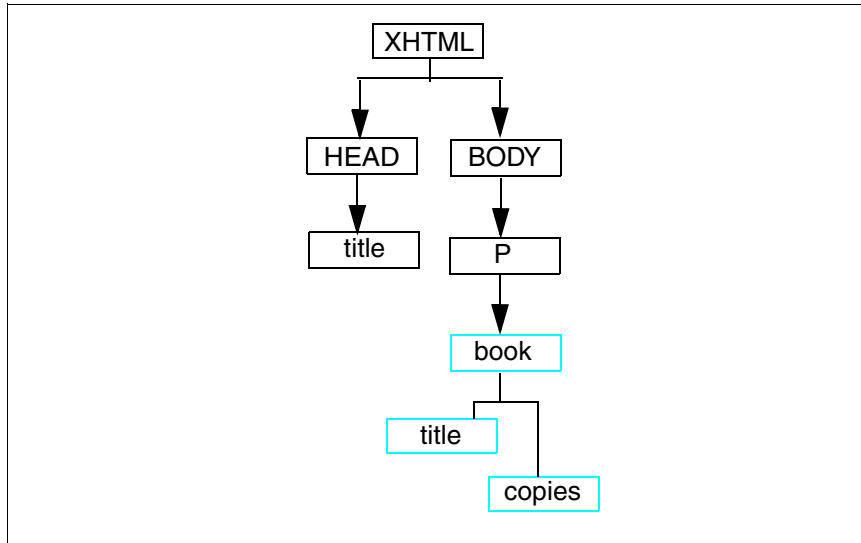


Figure 1-3 DOM tree after XSL transformation

Based on how we instruct the XSLT processor to access the source of the data being transformed, the processor will incrementally build the result by adding the filled-in templates. We write our stylesheets, or “transformation specifications,” primarily with declarative constructs, though we can employ procedural techniques if and when needed. We assert the desired behavior of the XSLT processor based on conditions found in our source.

Note that XSLT only manipulates the source tree and that the original XML document is left unchanged.

The most important aspect of XSLT is that it allows you to perform extremely complex manipulations on the selected tree nodes by affecting both content and appearance. Indeed, the final output may bear absolutely no resemblance to the source document. This ability to manipulate the nodes is where XSLT far surpasses CSS.

The World Wide Web Consortium (W3C) has set the recommended standards for XSLT Version 1.0. The W3C proposed recommendation for XSL is available at the following URL:

<http://www.w3.org/TR/xslt>

1.2.5 XPath

XPath is a string syntax for building addresses to the information found in an XML document. We use this language to specify the locations of document structures or data found in an XML document when processing that information using XSLT. XPath allows us from any location to address any other location or content. In other words, XPath is a tool used in XSLT to select certain information to be formatted.

XPath 2.0 is the current version of the language. A number of implementations exist but are not as widely used as XPath 1.0. The XPath 2.0 language specification changes some of the fundamental concepts of the language such as the type system.

For details, see <http://www.w3.org/TR/xpath20/>

XPath expressions are usually built out of patterns, which describe a branch of an XML tree. A pattern therefore is used to reference one or more hierarchical nodes in a document tree.

Some XPath patterns are shown in Table 1-1. These are just a few examples to give you an idea what kind of things can be selected.

Table 1-1 XPath

Symbol	Meaning
/	Root pattern. It refers to immediate child.
//	Separator of steps. It refers any descendant in the the node. By default of axis ^a , in the next step it refers to a child.
.	Context item. It refers to current node.
*	Wildcard pattern. It refers to all elements in the actual node.
@	It refers to an attribute by preceding an attribute name.
@*	Refer to all attributes in the actual node.

a. An axis defines a node-set relative to the current node.

XPath models an XML document as a tree of nodes, as follows:

- ▶ A root node
- ▶ Element nodes
- ▶ Attribute nodes
- ▶ Text nodes
- ▶ Namespace nodes
- ▶ Processing instruction nodes
- ▶ Comment nodes

The basic syntactic construct in XPath is the *expression* (see Example 1-10). An object is obtained by evaluating an expression, which has one of the following four basic types:

- ▶ Node-set (an unordered collection of nodes without duplicates)
- ▶ Boolean
- ▶ Number
- ▶ String

Example 1-10 XPath

```
<?xml version="1.0"?>
<!DOCTYPE library system "library.dtd">
<library>
  <book ID="B1.1">
    <title>xml</title>
    <copies>5</copies>
  </book>
  <book ID="B2.1">
    <title>WebSphere</title>
    <copies>10</copies>
  </book>
  <book ID="B3.2">
    <title>great novel</title>
    <copies>10</copies>
  </book>
  <book ID="B5.5">
    <title>good story</title>
```

```

    <copies>10</copies>
  </book>
</library>

```

Considering Example 1-10, we could make paths like:

- ▶ `/book/copies`
Selects all `copies` element children of `book` elements.
- ▶ `/book//title`
Selects all `title` elements in the tree, although `title` elements are not immediate children.
- ▶ `/book/@ID`
Selects all `ID` attributes for `book` elements.

But as we mentioned previously, it is also possible to select elements based on other criteria, such as:

- ▶ `/library/*/book[title eq "good story"]`
Selects all `book` elements beyond `library` element, but only if the `title` element matches with `good story`.

1.2.6 XQuery

XQuery is a functional language that extends XPath. Its basic building blocks are expressions constructed from keywords, operators (symbols), and operands (that are usually other expressions). Expressions can be nested with full generality. An XQuery query is composed of a prologue and a body. The query prologue is optional and consists of declarations that define the execution environment of the query. The query body consists of an expression that provides the result of the query. The input and the output of the query are values (instances) of the XQuery 1.0 and XPath 2.0 Data Model (XDM)⁴.

In Figure 1-11 we show a typical XQuery query. It begins with the `XQUERY` keyword followed with a prologue (optional) and a body. The prologue in our example contains default namespace declaration (the second line). The rest of the query is its body. It consists of one or more XQuery expressions.

Example 1-11 Sample XQuery

```

declare default element namespace "http://sample.name.space.com";
for $cust in db2-fn:xmlcolumn('XPS.DOC')
return
$cust/Name/LastName;

```

1.2.7 XHTML

The history of XHTML is very simple: it is derived directly from HTML version 4.01 and is designed to be used with XML. Indeed, XHTML is part of a whole new suite of “X” technologies, with acronyms such as XML, XPATH, XSL, and XSLT, that are destined to have a profound effect on the Internet.

⁴ The term XDM instance is used, like the term value, to denote an unconstrained sequence of nodes or atomic values in the data model.

People often think XML is an extension of HTML, but XHTML is the real extension of HTML.

There are a few fundamental differences between HTML and XHTML that significantly affect how you code with XHTML. While HTML is a loose and forgiving language, XHTML demands firm adherence to the rules of grammar.

Fortunately, the syntax and coding rules are very straightforward, easy to implement, and they make sense. The real purpose of these rules is to allow a seamless integration of XHTML with XML and other related X technologies. The rules are summarized as follows:

- ▶ All attributes, events, and tags must be written in lower case.
- ▶ All elements must be closed.
- ▶ The value assigned to an attribute must be enclosed in quotes.
- ▶ No attribute may be minimized.
- ▶ All elements must be properly nested.
- ▶ XHTML documents must be well-formed.
- ▶ There must be a DOCTYPE declaration.

Notice that this last rule implies that there must be a DTD to validate the XHTML document. HTML has become an XML document.

XHTML document types

XHTML 1.0 specifies three XML document types that correspond to three DTDs: Strict, Transitional, and Frameset. The most common is XHTML transitional. The DOCTYPE declaration at the beginning of the XHTML document specifies which type is being used.

▶ XHTML 1.0 Strict

Use this when you want really clean markup, free of presentational clutter. Use this together with Cascading Style Sheets. Example 1-12 shows a strict DTD.

Example 1-12 Strict DTD

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
```

▶ XHTML 1.0 Transitional

Use this when you need to take advantage of HTML's presentational features and when you want to support browsers that don't understand Cascading Style Sheets. Example 1-13 shows a transitional DTD.

Example 1-13 Transitional DTD

```
<!DOCTYPE html
PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd">
```

▶ XHTML 1.0 Frameset

Use this when you want to use HTML Frames to partition the browser window into two or more frames. Example 1-14 shows a frameset DTD.

Example 1-14 Frameset DTD

```
<!DOCTYPE html
PUBLIC "-//W3C//DTD XHTML 1.0 Frameset//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-frameset.dtd">
```

XLink

XML Linking Language (XLink) is a powerful and compact specification for the use of links in XML documents.

Every developer is familiar with the linking capabilities of the Web today. However, as the use of XML grows, we quickly realize that simple tags like the following ones in Example 1-15 are not going to be enough in a near future.

Example 1-15 Simple tags

```
<a href="elem_lessons.html">Freud</a information about X > are not going to be  
enough for many of our needs.
```

XLink allows elements to be inserted into XML documents to create and describe links between resources. It uses XML syntax to create structures that can describe links similar to the simple unidirectional hyperlinks of today's HTML, as well as more sophisticated links.

XLink provides a framework for creating both basic unidirectional links and more complex linking structures. It allows XML documents to:

- ▶ Assert linking relationships among more than two resources
- ▶ Associate metadata with a link
- ▶ Express links that reside in a location separate from the linked resources

Even though XLink has not been implemented in any of the major commercial browsers yet, its impact will be crucial for the XML applications of the near future. Its extensible and easy-to-learn design should prove an advantage as the new generation of XML applications develop.

For more information about Xlink, refer to the specification document from W3C:

<http://www.w3.org/TR/xlink/>

XPointer

XML Pointer Language (XPointer) specifies a language that builds upon the XPath, to support addressing into the internal structures of XML documents. In particular, it provides for specific references to elements, character strings, selections, and other parts of XML documents—whether or not they bear an explicit ID attribute—using traversals of a document's structure and choice of parts based on their properties, such as element types, attribute values, character content, and relative position, containment, and order. Xpointer defines the meaning of the “selector” or “fragment identifier” portion of URIs that locate resources of MIME media types text/xml and application/xml.

In XPointer, one defines the addressing expression to link XML documents using XPath. For more information about XPointer, refer to the specification documents from W3C:

<http://www.w3.org/TR/xptr/>

1.2.8 XSL, XSLT, Xpath, and XHTML examples

Let's first look at some example stylesheets using two implementations of XSLT 1.0 and XPath 1.0

Consider the XML file shown in Example 1-16. It is a very simple file we are going to use as the source of information for our XSLT transformation.

Example 1-16 hello.xml

```
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="hello.xsl"?>
<greeting>Hello world.</greeting>
```

Note that the stylesheet association processing instruction in line 2 refers to a stylesheet with the name `hello.xsl` of type XSL. Recall that an XSLT processor is not obliged to respect the stylesheet association preference, so let us first use a standalone XSLT processor with the stylesheet `hellohtm.xsl`, shown in Example 1-17.

Example 1-17 hellohtm.xsl

```
<?xml version="1.0"?><!--hellohtm.xsl-->
<html xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
xsl:version="1.0">
<head><title>Greeting</title></head>
<body><p>Words of greeting:<br/>
<b><i><u><xsl:value-of select="greeting"/></u></i></b>
</p></body>
</html>
```

This file looks like a simple XHTML file: an XML file using the HTML vocabulary. Indeed, it is just that, but we are allowed to inject into the instance XSLT instructions using the prefix for the XSLT vocabulary declared in line 3. We can use any XML file as an XSLT stylesheet provided it declares the XSLT vocabulary within and indicates the version of XSLT being used. Any prefix can be used for XSLT instructions, though convention often sees XSL: as the prefix value.

The `xsl:value-of` instruction uses an XPath expression in the `select=` attribute to calculate a string value from our source information. XPath views the source hierarchy using parent/child relationships. The XSLT processor's initial focus is the root of the document, which is considered the parent of the document element. Our XPath expression value `"greeting"` selects the child named `"greeting"` from the current focus, thus returning the value of the document element named `"greeting"` from the instance.

We invoke the XSLT processor to point to which is the XML source file, which is the XSL stylesheet, and where to leave the result. Example 1-18 shows the result file.

Example 1-18 Output from XSLT processor

```
<html>
<head>
<title>Greeting</title>
</head>
<body>
<p>Words of greeting:<br>
<b><i><u>Hello world.</u></i></b>
</p>
</body>
</html>
```

This is an HTML file, any browser could interpret it as shown in Example 1-4.

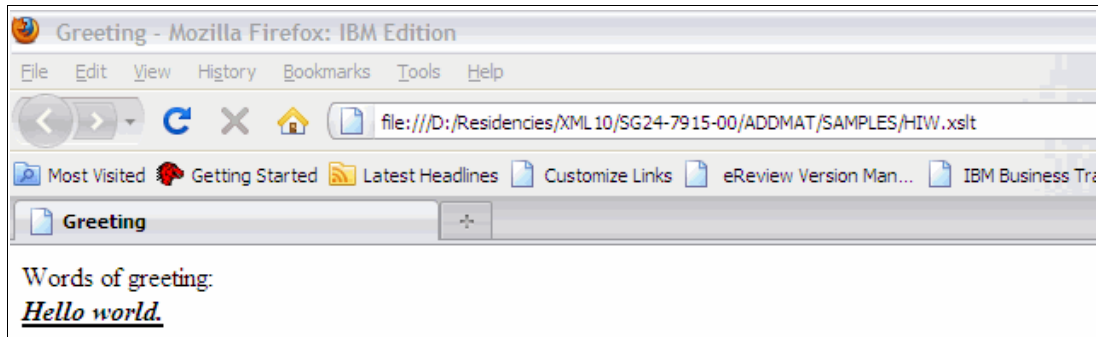


Figure 1-4 Hello world

1.3 What is in this book

In this book we provide information in four areas related to the use of pureXML in a DB2 for z/OS environment:

- ▶ Level set information on XML

We describe XML basics in 1.1, "Importance of XML data" on page 2 and 1.2, "XML introduction" on page 4 of this chapter.

Chapter 2, "XML and DB2 for z/OS" on page 21 starts describing the pureXML support in DB2 for z/OS.

- ▶ Introducing XML in a DB2 for z/OS environment

Chapter 3, "Application scenario" on page 45 describes the scenario which we will use across the various implementation steps.

Chapter 4, "Creating and adding XML data" on page 51 explains how to define XML objects in a DB2 for z/OS environment.

Chapter 5, "Validating XML data" on page 73 introduces the need for validating documents and the techniques to do so.

Chapter 6, "DB2 SQL/XML programming" on page 87 gives the basis on SQL/XML programming.

- ▶ Application development

We document the steps for the implementation of a simple but meaningful XML application scenario. We have chosen to provide samples in COBOL and Java language. The purpose being to provide an easy path to follow to integrate the XML data type for the traditional COBOL/DB2 user or the more innovative Java developer.

Chapter 7, "Using XML with Java" on page 129 describes the Java implementation.

Chapter 8, "Using XML with COBOL" on page 155 describes the COBOL implementation.

- ▶ Database administration

We also add considerations for the data administrator and suggest best practices for ease of use and better performance.

Chapter 9, "Utilities with XML" on page 181 revisits most of the DB2 utilities when used for XML data type.

Chapter 10, "XML-related tasks for the DBA" on page 229 highlights differences in administering XML data.

Chapter 11, “Performance considerations” on page 243 provides a checklist of the major performance considerations when deploying an application that uses pureXML.



XML and DB2 for z/OS

In this chapter we provide a concise overview of the XML capabilities within DB2. The intent of this chapter is to introduce the major elements of pureXML, and to explain how DB2's support for XML data makes it very easy to store and process XML data in an efficient and productive way.

We also summarize the XML infrastructure within DB2 that is required to support these XML capabilities. Most of the XML function is available “out of the box” after a standard DB2 10 for z/OS installation, but some facilities (such as XML schema validation) require that optional parts of the installation process are completed. We clarify these optional steps.

If you want a brief introduction to DB2 pureXML without excessive technical detail, you should read this chapter. This chapter also acts as a level setter for those who want to understand the technical details of pureXML which is covered in far more detail in the subsequent chapters of this book.

This chapter contains the following sections:

- ▶ XML capabilities provided by DB2
- ▶ Supporting infrastructure
- ▶ Choice of tools

2.1 XML capabilities provided by DB2

DB2 allows XML documents to be stored, managed and accessed as first class objects within a DB2 database. XML documents can contain large and flexible data structures. Not only can these data structures be stored in DB2 (which many relational databases allow to some extent), but additionally the nodes and elements within an XML document can be accessed and indexed to the finest level of granularity.

It is the “native” support for XML that makes DB2 pureXML so powerful, because XML documents and schemas can be used within DB2 with minimal administration work, as we will show within the application scenario that is used by this book.

Some of the main capabilities provided by DB2 pureXML are:

- ▶ A native XML data type.
- ▶ SQL/XML language, providing XML functions within the SQL language to access XML structures with the full power of XPath expressions.
- ▶ Hybrid data access, whereby relational and XML structures can be accessed together using a single SQL/XML statement.
- ▶ Read and write access to XML documents and sub-documents.
- ▶ XML Indexes (based on XPath expressions), to provide efficient access paths.
- ▶ XML schema validation (against an XML schema registered in the Schema Repository) including support for multiple versions of XML schemas.

Collectively, these capabilities allow you to use DB2 as a repository for both relational and XML data structures. Starting with version 9, DB2 for z/OS is a hybrid database. You can store both kinds of structures in a single database, write applications that use both kinds of structures concurrently, and manage all your data with the same set of Database Administration utilities.

The XML model of data is significantly different from the relational model of data. Java and JDBC provide many powerful XML manipulation capabilities.

COBOL and PL/I have also added many XML manipulation capabilities, as described at:

<http://www.ibm.com/support/docview.wss?uid=swg27004198&aid=1>
<http://www.ibm.com/software/awdtools/cobol/zos/>

It is the range of XML capabilities provided by DB2, summarized above, that makes it very easy to incorporate XML data into all DB2 databases and applications, including those written in languages like COBOL and PL/I.

This section reviews the major XML capabilities of DB2, and explains why each of these capabilities is important for building DB2 applications that contain XML data.

We examine the following topics:

- ▶ Native XML data type
- ▶ SQL/XML language
- ▶ Hybrid data access
- ▶ XML update
- ▶ XML indexes
- ▶ XML schema repository and schema validation

2.1.1 Native XML data type

All data types are equal, but some are more equal than others. To think of XML as a data type, equivalent in scope and importance to an integer for example, is misleading. XML is implemented as a data type within DB2, but it also encompasses an entire data model that contains many other data types (including the aforementioned integer).

Given the proliferation of XML within modern systems, it is imperative that a database should be able to store XML documents efficiently, and support efficient data access to any data element within the XML documents (alongside all the other qualities of services that a database must provide).

Some vendors have provided XML-only databases to do this. Other relational database vendors have provided ways of storing XML documents, and “stripping” out important data elements into relational structures (like DB2 V7 and V8 did).

The XML data type is part of the ANSI SQL standard. DB2 implements this standard, and additionally provides constructs such as XML indexes to make it productive and performant.

What DB2 pureXML provides is the combination of native XML data support alongside its established relational data support, in a fully integrated way. The XML data type (first introduced in DB2 9) is the most fundamental component of DB2’s support for XML. It is very different from storing XML as a string data type (as was done in DB2 V7 and V8).

Before the native XML data type existed

When XML was stored as a string (in DB2 V8 and earlier), DB2 had no inherent understanding of the structure within the XML document. DB2 V7 provided a facility called the XML Extender, which helped you to parse an XML document which was stored in DB2 as a string, and optionally strip out XML elements of interest. However, you had to perform this parsing every single time you wanted to access an XML document. If you wanted to process the data within the stored XML documents, you had two choices:

- ▶ You could either read and parse the entire document in order to access the data elements within the XML documents. (which could be very expensive, particularly if you were processing many rows through a cursor operation).
- ▶ Or you could perform a lot of up-front database administration and development work to strip out the data elements that might be used for searching and joining, and store them in additional DB2 columns with traditional DB2 data types.

The first approach was not practically viable from a performance or cpu-cost point of view. XML parsing is a CPU-intensive activity which you do not want to incur every single time you access the XML data.

The second approach is practical, but is very unproductive. This approach requires a significant development project to be undertaken to prepare the DB2 database before you can start to develop the application that will use the XML data.

After the native XML data type was introduced

The native XML data type (provided by DB2 9 and further improved by DB2 10) allows XML documents to be stored within DB2 in a very practical and productive way.

- ▶ Performance: The XML document is parsed once only, as it is inserted (or loaded) into DB2. The internal structure of the XML document is then accessible (without being re-parsed) by the SQL/XML functions provided by DB2.
- ▶ Productivity: No stripping of XML elements into separate DB2 columns is required (although if you choose to store some XML data elements in DB2 columns for database

design reasons, it is very easy to do this). All the XML elements can be referenced and indexed where they reside within the XML column.

XML documents are placed inside DB2 by inserting them into a column that has been defined with the XML data type, as illustrated in Example 2-1.

Example 2-1 Defining and populating an XML column

```
CREATE TABLE XMLR3.XMLTEST_TAB (  
    TESTKEY BIGINT,  
    TESTXMLCOL XML )  
    IN XMLR3DB.TEST_TS ;  
  
INSERT INTO XMLR3.XMLTEST_TAB ( TESTKEY, TESTXMLCOL )  
    VALUES ( 1, '<customerinfo><name>Amir Malik</name>  
    <phone type="work">408-555-1358</phone></customerinfo>' ) ;
```

The next sections describe how XML data can be processed once it is inside a a DB2 table, stored in a column of the XML data type.

2.1.2 SQL/XML language

Now that you can store XML as a native data type within DB2, you need a data access language to work with it.

Plain old SQL is able to access tables with XML columns, but it does not have the direct manipulation capabilities to do anything meaningful with the XML structure within the retrieved XML document. That is why DB2 has implemented SQL/XML extensions to SQL, to provide a range of functions that allow the contents of XML documents to be processed directly. These functions can also be encapsulated in DB2 stored procedures and user defined functions, so that 'plain old SQL' can be used against XML data with functions and procedures that were developed using SQL/XML extensions.

SQL/XML is an extension to the SQL standard as defined by ISO/IEC 9075-14:2003. The SQL/XML part of DB2 pureXML, being standards-based, is supported by many other databases. It provides a range of XML related extensions to the SQL language, to allow XML data to be accessed. The functions provided by SQL/XML can be categorized into three main groups.

- ▶ XML publishing functions, which allow XML documents to be created from the contents of relational data.
- ▶ XML handling functions which allow the user to embed XPath expressions in SQL statements. XPath expressions are a subset of the XQuery standard.
- ▶ XML conversion functions, which support the interchange of data between relational and XML models of data.

XML publishing functions

SQL/XML provides a range of "XML Publishing Functions" that allow XML documents to be created from the contents of relational data. Example 2-2 shows a simple example of XML Publishing Functions that conveys the concepts of what is possible. In this example we create a normal relational table and populate it. Then we generate an XML document based on the contents of the relational table.

Example 2-2 XML Publishing Functions of SQL/XML

```

Create Table Address (
  CUSTNAME  VARCHAR(70),
  STRTNM    VARCHAR(70),
  BLDGNB    VARCHAR(16),
  PSTCD     VARCHAR(16),
  TWNNM     VARCHAR(35)
) ;

Insert into Address
  values ('John Smith', 'Bailey Avenue', '555', '95141', 'San Jose') ;

SELECT
  XMLELEMENT(
    NAME "MsgRcpt",
    XMLELEMENT(NAME "Nm", CUSTNAME),
    XMLELEMENT(NAME "Pst1Adr",
      XMLELEMENT(NAME "StrtNm", STRTNM),
      XMLELEMENT(NAME "BldgNb", BLDGNB),
      XMLELEMENT(NAME "PstCd", PSTCD),
      XMLELEMENT(NAME "TwnNm", TWNNM)
    ) ) from Address ;

-- yields result

<MsgRcpt>
  <Nm>John Smith</Nm>
  <Pst1Adr>
    <StrtNm>Bailey Avenue</StrtNm>
    <BldgNb>555</BldgNb>
    <PstCd>95141</PstCd>
    <TwnNm>San Jose</TwnNm>
  </Pst1Adr>
</MsgRcpt>

```

XML Publishing Functions are very straightforward. They provide a range of string concatenation steps that can be combined together to form an XML documents. They simplify what you could already have done with the standard SQL CONCAT function.

XML handling functions

The XML handling functions are what enables SQL statements to access and manipulate the contents of XML documents stored within XML documents held within DB2. They include XMLQUERY, XMLTABLE and XMLEXISTS functions. Each of these functions makes use of XPath expressions, which is a subset of XQuery.

Existing DB2 programmers only need to extend their skills to include XPath expressions, in order to extend the relational programs that they already write, to incorporate XML data as well. Their capabilities are best explained with some simple examples, based on a simple table with two columns, illustrated in Example 2-3.

Example 2-3 Example table with XML column

```

Create Table XMLADDRESS (
  XID INT,

```

```
XMLADDRESS XML ) ;
```

```
Insert into XMLADDRESS ( XID , XMLADDRESS )
  values (1, '<MsgRcpt><Nm>John Smith</Nm><Pst1Adr>
    <StrtNm>Bailey Avenue</StrtNm><BldgNb>555</BldgNb>
    <PstCd>95141</PstCd><TwnNm>San Jose</TwnNm></Pst1Adr></MsgRcpt>' ) ;
```

```
Insert into XMLADDRESS ( XID , XMLADDRESS )
  values (2, '<MsgRcpt><Nm>John Doe</Nm><Pst1Adr>
    <StrtNm>Antelope Street</StrtNm><BldgNb>32</BldgNb>
    <PstCd>87233</PstCd><TwnNm>San Diego</TwnNm></Pst1Adr></MsgRcpt>' ) ;
```

XMLEXISTS

The XMLEXISTS function is used as a predicate to retrieve DB2 rows based on predicates that are applied to the data within an XML column, as illustrated in Example 2-4. In this example, each XML document is examined to see whether a data element <Nm> exists with a value of “John Doe” at XPath location /MsgRcpt, and returns the rows where the predicate is satisfied.

Example 2-4 Simple XMLEXISTS example

```
select c.xid, c.xmladdress
  from xmladdress c
  where xmlexists('$i/MsgRcpt[Nm = "John Doe"]'
    passing c.xmladdress as "i");
```

--yields result

XID XMLADDRESS

--- -----

2 <MsgRcpt><Nm>John Doe</Nm><Pst1Adr><StrtNm>Antelope...

Figure 2-1 may help to visualize the query.

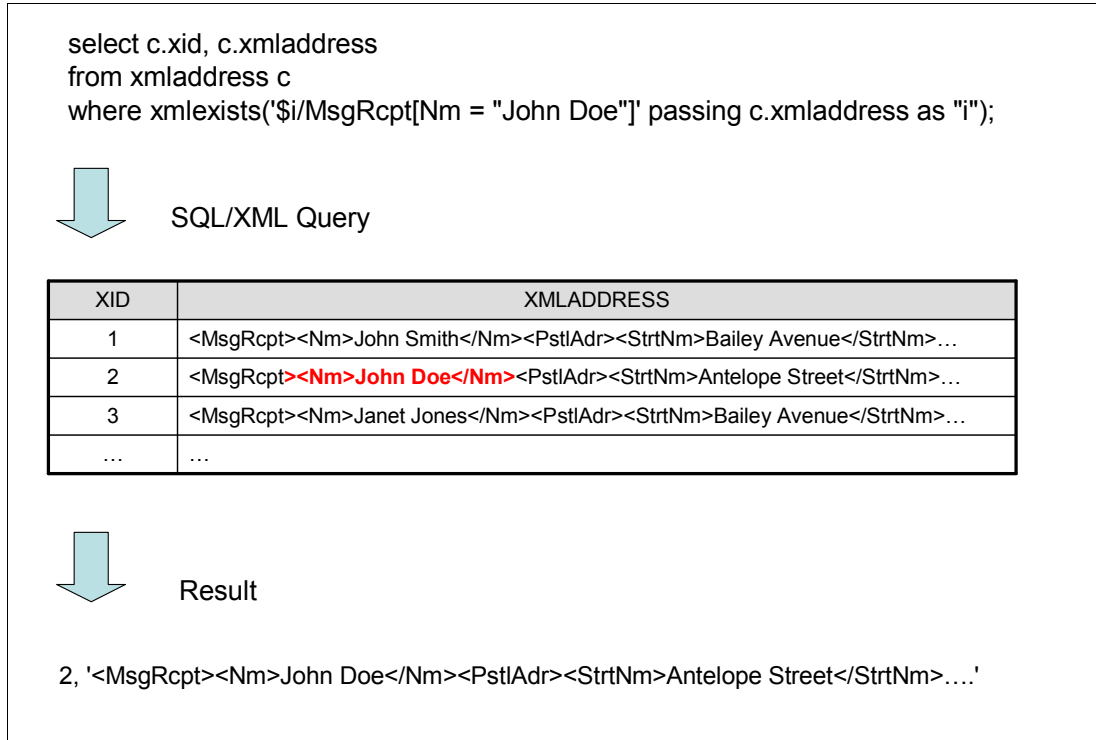


Figure 2-1 SQL/XML query with XMLEXISTS predicate

Let us take a moment to examine the structure of the XMLEXISTS predicate as the first example in this book of the SQL/XML extensions to the SQL language.

- ▶ The XMLEXISTS clause is the same as any other relational predicate, except that the XMLEXISTS clause contains XML syntax expressions.
- ▶ If the Nm field had been a VARCHAR column, the expression would have been where Nm = "John Doe"
- ▶ The XPath expression is applied to the XML column using the *passing* clause. In this way, the XML document in each row that is accessed is passed as "i" to the XPath predicate for evaluation
- ▶ The comparison that is made by the XMLEXISTS clause is identical to what would have been done relationally. The contents of the XML location /MsgRcpt/Nm are compared to the string value "John Doe".
- ▶ If the values match, the predicate is satisfied and the row qualifies
- ▶ If the values do not match, the row does not qualify
- ▶ If the data types do not match, the row does not qualify

Note that the SQL/XML statement does not know whether the contents at /MsgRcpt/Nm is a string or a number or any other data type. It determines this at runtime. The data type of the /MsgRcpt/Nm location may be constrained to be a particular type if the document conforms to an XML schema. XML schemas and validation of XML documents against XML schemas is covered shortly in 2.1.6, "XML schema repository and schema validation" on page 36.

Also, be aware of namespaces. Normally XML document will have a declared namespace. A namespace provides the ability to uniquely define a data element or attribute within an XML document, so that when a tag like <phone> is used more than once in an XML document, the precise context of that <phone> tag is understood.

If an XML document has a namespace declaration, then the query must also define the namespace, in order to ensure that we are referencing the correct XML data. Example 2-5 shows how a namespace would be declared on the XMLEXISTS function, if a namespace was required.

Example 2-5 XMLEXISTS example with namespace declaration

```
select c.xid, c.xmladdress
  from xmladdress c
  where xmlexists('
    declare default element namespace "http://www.names.com";
    $i/MsgRcpt[Nm = "John Doe"]'
    passing c.xmladdress as "i");
```

If we were to code the statement:

```
select * from table where name = 'John Doe'
```

the DBA or developer should be considering whether or not the name column should be indexed in order to avoid a table space scan.

Exactly the same consideration applies to XML data. If we really wanted to code the SQL/XML statement in Example 2-4, then we should be considering an XML index to provide an efficient access path to the <Nm> elements in the XML documents. Otherwise we will end up executing the XPath expression against every single XML document in the table.

Of course, the DB2 optimizer is able to choose an access path based on a combination of XML and relational predicates. So relational and XML indexes are both available for access path selection, and can both be used in the same access plan.

XMLTABLE

The XMLTABLE function is used to retrieve XML elements and attributes from an XML document and map them to a relational table structure, to be used by programs exactly as if the data had been retrieved from a wholly relational table, as illustrated in Example 2-6.

Example 2-6 Simple XMLTABLE example

```
SELECT X.NAME, X.STREET, X.STREETNUM, X.POSTCODE, X.TOWN
  FROM XMLADDRESS C,
  XMLTable('$cu/MsgRcpt/Pst1Adr'
    PASSING C.XMLADDRESS as "cu"
    COLUMNS
      "NAME"          CHAR(18)  PATH './Nm',
      "STREET"        CHAR(18)  PATH 'StrtNm',
      "STREETNUM"     CHAR(16)  PATH 'BldgNb',
      "POSTCODE"      CHAR(16)  PATH 'PstCd',
      "TOWN"          CHAR(18)  PATH 'TwnNm'
  ) AS X
```

... yields

NAME	STREET	STREETNUM	POSTCODE	TOWN
John Smith	Bailey Avenue	555	95141	San Jose
John Doe	Antelope Street	32	87233	San Diego

The XMLTABLE in Example 2-6 shows how individual elements within an XML document can be mapped to a relational structure, and retrieved alongside other relational columns from the DB2 table which the XML column belongs to. Figure 2-2 provides a graphical representation of the XMLTABLE function.

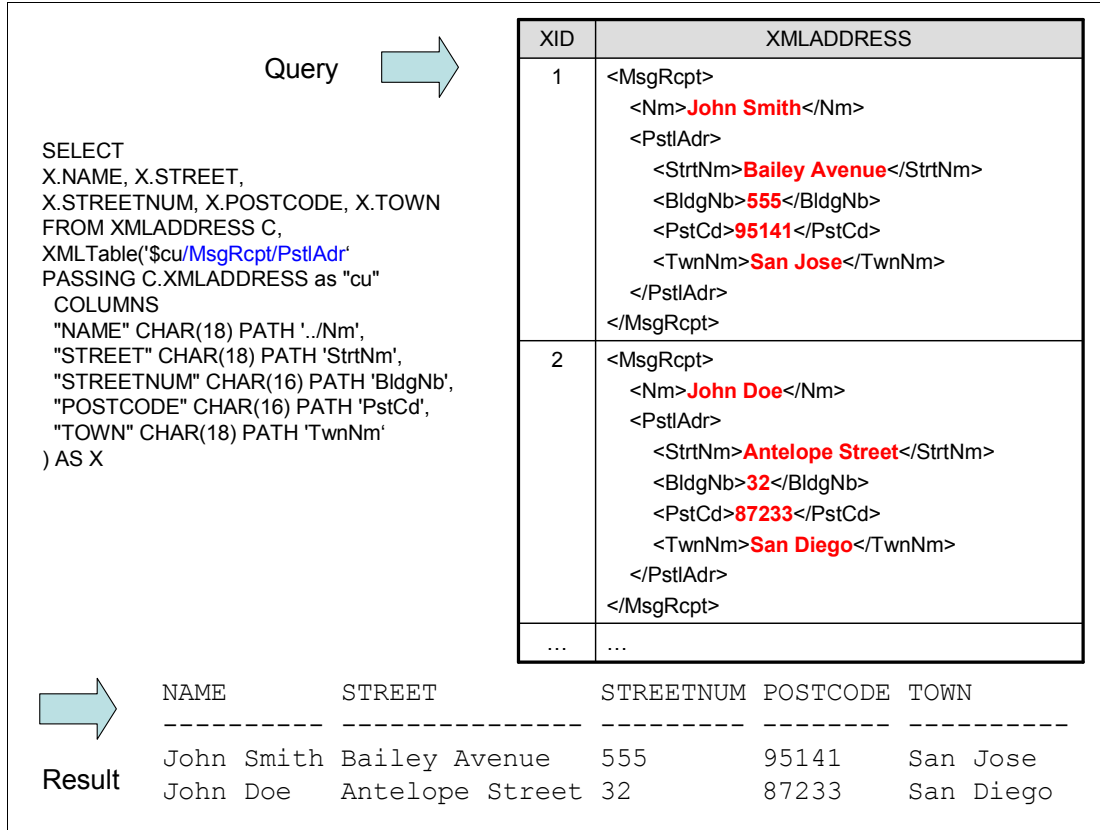


Figure 2-2 XMLTABLE function example

The mapping is based on an XPath expression which is used to navigate to a particular anchor point in the XML document (/MsgRcpt/PstlAdr) and then using relative XPath expressions (such as /StrtNm at the next level down, and ../Nm at the next level up from that anchor point to access XML data elements that are mapped to relational fields (such as STREET VARCHAR(18)).

Note that the example in Figure 2-2 has mapped data from Nm, StrtNm and TwnNm as CHAR(18), whereas the original source of the data for all three fields was VARCHAR(70) in Example 2-2 on page 25. The XMLTABLE function will perform a range of standard data type casting, and will return an error if the casting is not possible. For example, if you attempted to cast Nm as an integer, you would get SQL16061N. The value "John Smith" cannot be constructed as, or cast (using an implicit or explicit cast) to the data type "xs:integer".

More details will be covered later in the book on managing the XML to relational mapping and handling errors that might arise (for example, when casting an XML element to a specific relational data type).

XMLQUERY

The XMLQUERY function is used to embed XPath expressions within an SQL/XML statement. It always produces a column of type XML, and, as a scalar function, it returns a sequence of items for each document (row). This is illustrated in Example 2-7.

Example 2-7 Simple XMLQUERY example

```

select
  xmlquery('$i/MsgRcpt/Nm ' passing c.xmladdress as "i")
  as Names
  from xmladdress c;

-- yields

Names
-----
<Nm>John Smith</Nm>
<Nm>John Doe</Nm>

```

The XMLQUERY in Example 2-7 shows how XPath expressions can be executed within an SQL statement, with the resulting XML document (or subdocument) being returned as an XML structure. This particular example retrieves the XML structure at node /MsgRcpt/Nm for every single XML document within the XMLADDRESS table.

The XMLEXISTS and XMLTABLE functions are particularly attractive when using traditional programming languages, because they return data in a relational format. XMLTABLE can also return an XML column.

The XMLQUERY function is a little different because it returns a scalar value XML type, not necessarily an XML document (usually it is not a document). This does not necessarily introduce extra complexity to the traditional programmer, because the returned XML type can be cast to a string data type, or moved on to another repository like WebSphere Message Queue (MQ).

Collectively, the XMLEXISTS, XMLQUERY and XMLTABLE functions allow the traditional mainframe programmer to access and manipulate XML documents with a minimal increase in skills required (namely, the ability to write XQuery and XPath expressions, and embed them in SQL statements). We examine the considerations for using each of these functions with index access in Chapter 11, "Performance considerations" on page 243.

XML conversion functions

The XML conversion functions (XMLCAST, XMLPARSE and XMLSERIALIZE) support the interchange of data between relational and XML models of data.

XMLCAST will cast the contents of an XML data element to a relational data type as shown in Example 2-8.

Example 2-8 XMLCAST example casting a numeric data element to varchar or integer

```

select
  xmlcast(xmlquery('$i/MsgRcpt/Pst1Adr/BldgNb'
    passing c.xmladdress as "i") as varchar(10))
  as streetnumber from xmladdress c;

select
  xmlcast(xmlquery('$i/MsgRcpt/Pst1Adr/BldgNb'
    passing c.xmladdress as "i") as integer)
  as streetnumber from xmladdress c;

```

XMLPARSE parses a string of a well-formed XML document that conforms to XML 1.0 and

returns an XML type. XMLPARSE can be used in isolation or as part of an SQL INSERT operation as shown in Example 2-9. XMLPARSE also provides the option to strip whitespace, for storage efficiency, or preserve it.

Example 2-9 XMLPARSE function and whitespace handling

```
INSERT INTO XMLDEMO1 VALUES(201,
  XMLPARSE( DOCUMENT
    '<emails><email emailUse="work">      fred_smith@uk.ibm.com </email>
      </emails>')
  PRESERVE WHITESPACE ) );
```

```
INSERT INTO XMLDEMO1 VALUES(202,
  XMLPARSE( DOCUMENT
    '<emails><email emailUse="work">      fred_smith@uk.ibm.com </email>
      </emails>')
  STRIP WHITESPACE ) );
```

XMLSERIALIZE will convert an XML document into a serialized string value. Example 2-10 converts an XML document into a CLOB in UTF-8. The resulting data type could also be BLOB, DBCLOB, etc.

Example 2-10 XMLSERIALISZE example to convert an XML document to a UTF-8 CLOB

```
SELECT e.xid, XMLSERIALIZE(XMLADDRESS AS CLOB) AS "result"
  from xmladdress e;
```

2.1.3 Hybrid data access

One of the big advantages of adding XML support to DB2 is that hybrid database applications can be constructed, combining the relational model of data and the XML model of data.

A good example of a hybrid data application might be insurance quotes. When an internet user fills in multiple html forms to get an insurance quote, the insurance company wants to persist that information in a database, so that the internet user can come back at a later time to purchase the policy that was quoted for.

XML would be an excellent data model for storing insurance quotes because the data will typically contain a large number of data elements, and different quotes may be structured in many different ways. For example, the html forms filled in for a single driver car insurance policy for a small vehicle will differ significantly from the html forms filled in by a married couple insuring a powerful sports utility vehicle, and adding their 18 year old son as a named driver.

- ▶ If the quotation application was implemented in a relational structure, there may be 10 or 20 main tables, and multiple code tables.
- ▶ If the quotation application was implemented in an XML structure, a single XML document can contain all the structure and constraints for the quotations.

However, it is likely that the existing Client and Policy systems will have been written many years ago using the relational model of data. Hence, the new quotations application needs to be able to bridge the gap between XML and relational data models.

DB2's support for a hybrid data model makes it easy to develop a new internet quotes application. It is possible for the internet quotes application to store the quote as an XML document in a DB2 table. If the quote is subsequently taken up by the client, then the information stored within the XML document containing the quotation details can be retrieved using SQL/XML and re-used as input to the relational tables of the Client and Policy systems.

The XMLTABLE function provides the basis for a simple use case. The address details can be read from an XML Document containing the quotation, and inserted into an address table within the Policy system, as shown in Example 2-11.

Example 2-11 Hybrid data access example

```

INSERT INTO POLICY_SYSTEM.ADDRESS_TABLE ( ADDR1, TOWN, POSTCODE )
  SELECT X.ADDRESS, X.TOWN, X.POSTCODE
  FROM
    QUOTE_SYSTEM.XMLQUOTE_TABLE C,
    XMLTable('$cu/quote/addresses/address'
      PASSING C.CUSTCONTACTS as "cu"
      COLUMNS
        "ADDRESS" CHAR(20) PATH 'addressLine1',
        "TOWN" CHAR(20) PATH 'addressPostTown',
        "POSTCODE" CHAR(20) PATH 'addressPostCode'
    ) AS X
  WHERE X.QUOTE_ID = 123456 ;

```

Clearly this simplistic example does not reflect the way that Insurance systems may be architected, with error checking, input validation and so on. However, the example still serves to illustrate that the DB2 hybrid data model does provide an integrated database platform upon which it is easy to bridge the gap between relational and hierarchical models of data.

2.1.4 XML update

Full read and write access is provided for XML documents within DB2.

Inserting and updating XML data is supported with the XMLPARSE option, which will parse a string value to return an XML document, and optionally strip or preserve whitespace. The whitespace handling is illustrated in Example 2-9 on page 31.

Updating XML documents can be performed by either replacing the whole document, or updating a part of it. A full replacement of an XML document is coded with a simple SQL update statement, and an XMLPARSE function call if the source data is in string format. A partial document update uses the XMLMODIFY function to change an existing XML document.

The XMLMODIFY function depends on the underlying table space being a universal table space, so that multiple XML versions are supported. XML versioning is implemented in DB2 10 to improve concurrency and reduce locking: A new version of the XML document is created during an update, which allows SQL read operations (with the appropriate isolation level) to access the old version of the XML document concurrently.

The XMLMODIFY function operates on a node within an XML document, and has three usage variations: replace, insert or delete a node. We will show examples of all three variations of XMLMODIFY, based on the initial table contents shown in Example 2-12.

Example 2-12 Initial contents of XMLADDRESS table

```
XID XMLADDRESS
-----
 2 <MsgRcpt>
   <Nm>John Doe</Nm>
   <Pst1Adr>
     <StrtNm>Antelope Street</StrtNm>
     <BldgNb>32</BldgNb>
     <PstCd>87233</PstCd>
     <TwnNm>San Diego</TwnNm>
   </Pst1Adr>
 </MsgRcpt>
```

XMLMODIFY to replace a node

The first example replaces the node /MsgRcpt/Nm with another XML fragment. The replacement XML fragment is a literal string value, which is parsed using the XMLPARSE function.

Example 2-13 XMLMODIFY to replace a node

```
UPDATE XMLR3.XMLADDRESS C
  SET C.XMLADDRESS = XMLMODIFY(
    'replace node /MsgRcpt/Nm with $x',
    XMLPARSE('<Nm>Johnny Doe</Nm>') AS "x")
  where C.XID = 2
```

-- changes the second row in the XMLADDRESS table to the following:

```
XID XMLADDRESS
-----
 2 <MsgRcpt>
   <Nm>Johnny Doe</Nm>
   <Pst1Adr>
     <StrtNm>Antelope Street</StrtNm>
     <BldgNb>32</BldgNb>
     <PstCd>87233</PstCd>
     <TwnNm>San Diego</TwnNm>
   </Pst1Adr>
 </MsgRcpt>
```

XMLMODIFY to delete a node

The second example deletes the <Pst1Adr> node from the document.

Example 2-14 XMLMODIFY to delete a node

```
UPDATE XMLR3.XMLADDRESS C
  SET C.XMLADDRESS = XMLMODIFY(
    'delete node /MsgRcpt/Pst1Adr')
  WHERE C.XID = 2
```

-- changes the second row in the XMLADDRESS table to the following:

```
XID XMLADDRESS
-----
 2 <MsgRcpt>
```

```
<Nm>Johnny Doe</Nm>
</MsgRcpt>
```

XMLMODIFY to insert a node

The third example inserts the node <Notes> as the last child node at XPath location /MsgId.

Example 2-15 XMLMODIFY to insert a node

```
UPDATE XMLR3.XMLADDRESS C
  SET C.XMLADDRESS = XMLMODIFY(
    'insert node $x as last into /MsgRcpt',
    XMLPARSE('<Notes>Testing</Notes>') AS "x")
  where C.XID = 2
```

changes the second row in the XMLADDRESS table to the following:

```
XID XMLADDRESS
-----
 2 <MsgRcpt>
    <Nm>Johnny Doe</Nm>
    <Notes>Testing</Notes>
  </MsgRcpt>
```

In all cases, the following XML schema validation considerations apply:

- ▶ If the table has an XML type modifier for the INFO column, then the resultant XML value will be checked for well formedness and will also be checked for conformance to the registered XML schema.
- ▶ If the table does not have an XML type modifier for the INFO column, then the resultant XML value will be checked for well formedness only. XML schema validation could be requested manually, by calling the DSN_XMLVALIDATE function.

Sub-document update becomes important for performance reasons if small changes need to be made to large XML documents. For example, changing the details of one book in an XML document that contains a book catalog would be a lot cheaper than replacing and revalidating the whole document.

2.1.5 XML indexes

XML indexes are just as important to DB2 as relational indexes are. The best way to introduce XML indexes is to start with the similarities to relational indexes, and then to identify the differences.

The major reasons for defining relational indexes in DB2 are:

- ▶ Efficient access path to data (minimizing I/O and CPU resource consumption)
- ▶ Option to apply a unique constraint

These same reasons also apply to XML indexes. Their fundamental purposes are to facilitate performance and to enforce uniqueness.

However, the nature of an XML index has several differences from a relational index. The major differences are:

- ▶ Relational indexes may be defined on one or more relational columns. XML Indexes can only be defined on one XML element or XML attribute (using an XML pattern expression)

- ▶ Relational indexes always have one index entry for every row in a table. However XML indexes are much less prescriptive. XML indexes are based on an XML pattern. An XML pattern may occur any number of times in an XML document. So, XML indexes may contain 0, 1 or many entries for each row in the table.
- ▶ Relational indexes are always based on the data types of the column(s) that they are defined on. The data types found at the locations of an XML pattern may be many and varied unless an appropriate XML schema is enforced. Using XML schema validation against a well defined XML schema will allow the data types of XML elements to be controlled (as covered in 2.1.6, “XML schema repository and schema validation” on page 36).
- ▶ Relational indexes can be used to support table clustering. XML indexes may not be used for table clustering support.

XML indexes have a more complex set of physical design considerations than relational indexes. A good understanding of XPath expressions and XML schemas and XML namespaces is essential to designing XML indexes that are likely to be chosen by the DB2 optimizer.

Having sounded a caution that XML index design requires careful consideration of a number of new factors, it is well worth citing how powerful and productive they are. Without XML indexes, a database application would have to go through an extra development phase where key data elements are stripped to relational tables, and indexed using traditional relational indexes. With XML indexes, you can avoid this extra development phase and index your XML data every bit as efficiently as you index your relational data, and realize the simplicity that DB2’s hybrid data model provides.

Example 2-16 provides simple examples of how XML indexes are created in DB2.

Example 2-16 XML Index creation examples

```
CREATE INDEX XMLCITY ON XMLDEMO1(CUSTCONTACTS)
  GENERATE KEY USING XMLPATTERN
    '/employeeContacts/addresses/address/AddressPostCode'
  AS SQL VARCHAR(30);
```

In this example the contents of XPath location '/employeeContacts/addresses/address/AddressPostCode' are cast as a VARCHAR(30) data type, and an XML index is built.

- ▶ If the data type casting is successful, an index entry will be created.
- ▶ If the XPath does not exist, no index entry will be created.
- ▶ If the data type casting is unsuccessful
 - For numeric, date, and timestamp index data types, error-tolerance applies.
 - For VARCHAR indexes, where the failure of casting is due to length, either CREATE INDEX will fail or INSERT will fail.

An XML index could point to a small number of entries (or none) if the XML pattern of the index does not fit the data.

XML index design requires the nature of XML indexes to be understood, and the matching success of the index to be checked.

The DB2 optimizer can use XML indexes and traditional DB2 indexes together to produce the most efficient access paths for hybrid data access. The way that XML and relational indexes

work together is one of the best illustrations of how DB2 makes hybrid data both valuable and performant.

2.1.6 XML schema repository and schema validation

DB2 provides many features that allow the database administrator to implement and enforce a relational data model. Unique indexes, primary and foreign keys, check constraints, referential constraints etc.

DB2 also provides the ability to implement and enforce an XML data model, through its ability to validate XML documents against a registered XML schema. XML data structures can be defined via XML schemas. An XML Schema defines the structure of an XML document by defining:

- ▶ The elements that can appear in a document
- ▶ The attributes that can appear in a document
- ▶ Which elements are child elements
- ▶ The order of child elements
- ▶ The number of child elements
- ▶ Whether an element is empty or can include text
- ▶ The data types for elements and attributes
- ▶ The default and fixed values for elements and attributes

A very simple XML schema example is shown in Example 2-17 which defines how email information about an individual should be stored. The element <emails> can be a parent to between 1 and 5 specific <email> elements. Each email element is required to have an emailUse tag, which describes the nature of that particular email account (work, home etc...).

Example 2-17 Simple XML schema example

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">
<xs:element name="emails" maxOccurs="1" minOccurs="1">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="email" maxOccurs="5" minOccurs="1">
        <xs:complexType>
          <xs:simpleContent>
            <xs:extension base="xs:string">
              <xs:attribute name="emailUse" use="required">
                <xs:simpleType>
                  <xs:restriction base="xs:string">
                    <xs:enumeration value="home"/>
                    <xs:enumeration value="work"/>
                  </xs:restriction>
                </xs:simpleType>
              </xs:attribute>
            </xs:extension>
          </xs:simpleContent>
        </xs:complexType>
      </xs:element>
    </xs:sequence>
  </xs:complexType>
</xs:element>
</xs:schema>
```

The XML document in Example 2-18 conforms to the XML schema that is described in Example 2-17.

Example 2-18 XML document that conforms to previous XML schema

```
<emails>
  <email emailUse="work">UK000003@uk.ibm.com</email>
  <email emailUse="home">julie.woollacot@isp.com</email>
</emails>
```

DB2 support for XML schema validation consists of

- ▶ An XML schema repository, where XML schemas can be stored within DB2.
- ▶ A schema validation function, which allows an XML data type to be validated against a schema within the repository.
- ▶ A DDL option to define an XML column with a type modifier, so that schema validation is enforced by DB2 automatically for that column.

The schema repository is populated by using the four XML schema repository stored procedures that are provided by DB2. These stored procedures are invoked by commands from z/OS UNIX® System Services or DB2 for z/OS¹ and DB2 for Linux®, UNIX, and Windows® Command Line Processor (CLP), as shown in the following examples. The script shown in Example 2-19 shows a schema being registered to the DB2 XML schema repository, and then being completed.

Example 2-19 XML schema registration

```
register xmlschema http://www.mymodel
  from file://C:\SCHEMAPATH\myschema.xsd
  as SYSXSR.MYXMLSCHEMA ;

complete xmlschema SYSXSR.MYXMLSCHEMA ;
```

The SQL statement in Example 2-20 shows the DSN_XMLVALIDATE function being invoked manually, as part of an INSERT operation to a table (TABLE1).

Example 2-20 SYSXSR.DSN_XMLVALIDATE example

```
INSERT INTO TABLE1(XMLCOL1)
  VALUES (SYSIBM.DSN_XMLVALIDATE(:xml doc, 'SYSXSR.MYXMLSCHEMA')
);
```

The DDL statement in Example 2-21 shows another Table (TABLE2) being defined, where all XML documents stored in the XML column of that table will be automatically validated.

Example 2-21 XML type modifier example

```
CREATE TABLE XMLR3.TABLE2 (
  XMLCOL1 XML(XMLSCHEMA ID SYSXSR.MYXMLSCHEMA))
  IN XMLR3DB.XMLR3TS ;
```

¹ See DB2 for z/OS CLP at

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db29.doc.apsg/db2z_runsfromclp.htm

The DB2 XML schema validation services are essential for enforcing the integrity of XML documents against the rules that they should conform to. They are simple to use, and allow very complex XML schemas to be incorporated easily, as shown by the ISO20022 examples later in this book.

2.2 Supporting infrastructure

There is very little extra customization that must be done during a DB2 installation to take advantage of the pureXML capability within DB2. The XML data type is a standard DB2 data type that works *out of the box* with DB2.

XML validation is an optional facility requiring the XML Schema Repository (XSR). It is possible that you might not want to use DB2's XML schema validation for documents that have already been validated outside of DB2. However, in general, you need to have access to this facility or an equivalent one when you use DB2 pureXML.

For the most current maintenance level, see info APAR II14426. It contains a summary and pointers to all the XML support delivery APARs.

2.2.1 XSR installation steps

XSR has a pre-requisite that the following software is installed and configured.

- ▶ Workload Manager for z/OS (WLM)
- ▶ z/OS XML System Services
- ▶ Java 2 Technology Edition, V5 or later, 31-bit version
- ▶ IBM Data Server Driver for JDBC and SQLJ

XSR consists of four DB2-supplied stored procedures, one DB2 database, five table spaces, 8 tables and 13 indexes. Installation job DSNTIJRT is designed to create all DB-supplied stored procedures including these objects.

The four stored procedures, with a brief description of their purpose, are:

- ▶ SYSPROC.XSR_REGISTER is a C stored procedure, which registers an XML schema as a primary schema document in the DB2 XML schema repository.
- ▶ SYSPROC.XSR_ADDSCHEMADOC is a C stored procedure, which registers an additional XML schema document to an XML schema in the DB2 XML schema repository.
- ▶ SYSPROC.XSR_COMPLETE is a Java stored procedure, which completes the registration process of an XML schema.
- ▶ SYSPROC.XSR_REMOVE is a C stored procedure, which removes a registered XML schema from the XML schema repository.

The installation steps of the XML Schema Repository and functions are:

1. Customize and run jobs DSNTIJRW and DSNTIJMV

Define the WLM environment and startup procedure for the C language XML schema repository stored procedures. A dedicated WLM environment and startup procedure is required for the XSR stored procedures written in C (XSR_ADDSCHEMADOC, XSR_REGISTER, and XSR_REMOVE).

- Installation job DSNTIJRW installs and configures a WLM environment with the default name of DSNWLM_XML which you can use it to run the XML schema repository stored procedures.

- Installation job DSNTIJMV installs a WLM startup procedure named ssnmWLMX for that WLM environment.

Define the WLM environment and startup procedure for the Java language XML schema repository stored procedure. A dedicated WLM environment and startup procedure is required for the XSR stored procedure written in Java (XSR_COMPLETE).

- Installation job DSNTIJRW installs and configures a WLM environment with the default name of DSNWLM_JAVA which you can use it to run the XML schema repository stored procedures.
- Installation job DSNTIJMV installs a WLM startup procedure named ssnmWLMJ for that WLM environment.

2. Customize and run job DSNTIJRT

Run installation Job DSNTIJRT to create the XSR objects (database, table spaces, tables, indexes, stored procedures) and bind the DB2 XSR packages for the stored procedures.

3. Bind the Universal JDBC Driver

Bind the packages for the IBM Data Server Driver for JDBC and SQLJ (on USS and/or Windows)

2.2.2 XSR installation validation

DB2 installation job DSNTIJRV is designed as an installation verification program for all the DB2-supplied procedures, including those used by the XSR. Run it to validate that all services are operational. An abridged version of the output of DSNTIJRV is shown in Example 2-22, which shows a report of the testing of the XSR procedures and the ODBC procedures.

Example 2-22 DSNTIJRV installation verification job output

```

DSNT040I 06.27.58 DSNTRVYF ROUTINE VALIDATION SUMMARY
      STATUS      SCHEMA          SPECIFIC NAME
      --  -
      . . . . .
/  PASSED  SYSPROC          XSR_ADDSCHEMADOC
/  PASSED  SYSPROC          XSR_COMPLETE
/  PASSED  SYSPROC          XSR_REGISTER
/  PASSED  SYSPROC          XSR_REMOVE
      . . . . .
/  PASSED  SYSIBM          SQLCOLUMNS
/  PASSED  SYSIBM          SQLCOLPRIVILEGES
/  PASSED  SYSIBM          SQLFOREIGNKEYS
/  PASSED  SYSIBM          SQLFUNCTIONCOLS
/  PASSED  SYSIBM          SQLFUNCTIONS
/  PASSED  SYSIBM          SQLGETTYPEINFO
/  PASSED  SYSIBM          SQLPRIMARYKEYS
/  PASSED  SYSIBM          SQLPROCEDURECOLS
/  PASSED  SYSIBM          SQLPROCEDURES
/  PASSED  SYSIBM          SQLSPECIALCOLUMNNS
/  PASSED  SYSIBM          SQLSTATISTICS
/  PASSED  SYSIBM          SQLTABLEPRIVILEGES
/  PASSED  SYSIBM          SQLTABLES
/  PASSED  SYSIBM          SQLUDTS
      . . . . .
DSNT033I DSNTRVYF VALIDATION PROGRAM ENDED, RETURN CODE = 0

```

2.2.3 XSR setup troubleshooting

If you have difficulties using the XML schema repository, it may be helpful to check the following list of potential setup issues.

Check that both the WLM application environments for XSR are started and available, as shown in Example 2-23.

Example 2-23 z/OS console display WLM APPLENV status

```
D WLM,APPLENV=DSNWLMDBOB_XML

RESPONSE=SC63
IWM029I 14.43.07 WLM DISPLAY 072
APPLICATION ENVIRONMENT NAME STATE STATE DATA
DSNWLMDBOB_XML AVAILABLE
ATTRIBUTES: PROC=DBOBWLMX SUBSYSTEM TYPE: DB2

D WLM,APPLENV=DSNWLMDBOB_JAVA

RESPONSE=SC63
IWM029I 14.43.39 WLM DISPLAY 074
APPLICATION ENVIRONMENT NAME STATE STATE DATA
DSNWLMDBOB_JAVA AVAILABLE
ATTRIBUTES: PROC=DBOBWLMJ SUBSYSTEM TYPE: DB2
```

Check that the XSR tables are created with the SQL statement in Example 2-24.

Example 2-24 Check XSR tables exist

```
select name from sysibm.systables
  where creator = 'SYSIBM' and name like 'XSR%'

XSRANNOTATIONINFO
XSRCOMPONENT
XSROBJECTCOMPONENTS
XSROBJECTGRAMMAR
XSROBJECTHIERARCHIES
XSROBJECTPROPERTY
XSROBJECTS
XSRPROPERTY
```

Check that the XSR schema validation routines are created with the SQL statement in Example 2-25.

Example 2-25 Check XSR routines exist

```
select name from sysibm.sysroutine
  where schema= 'SYSPROC' and name like 'XSR%'

XSR_ADDSCHEMADOC
XSR_COMPLETE
```


XSR_REGISTER
XSR_REMOVE

Check that the Java environment is correctly setup, by creating two Java user defined functions that return the name and version number of the Java environment. Submit the following DDL in Example 2-26 to create a couple of Java stored procedures.

Example 2-26 Creation of Java stored procedures

```
CREATE FUNCTION SYSADM.JAVDRVV ()
  RETURNS VARCHAR(100)
  FENCED NO SQL
  LANGUAGE JAVA
  SPECIFIC JAVDRVV
  EXTERNAL NAME 'com.ibm.db2.jcc.DB2Version.getVersion'
  WLM ENVIRONMENT WLMJAVA
  NO EXTERNAL ACTION
  NO FINAL CALL
  PROGRAM TYPE SUB
  PARAMETER STYLE JAVA;

CREATE FUNCTION SYSADM.JAVDRVN ()
  RETURNS VARCHAR(100)
  FENCED NO SQL
  LANGUAGE JAVA
  SPECIFIC JAVDRVN
  EXTERNAL NAME 'com.ibm.db2.jcc.DB2Version.getDriverName'
  WLM ENVIRONMENT WLMJAVA
  NO EXTERNAL ACTION
  NO FINAL CALL
  PROGRAM TYPE SUB
  PARAMETER STYLE JAVA;
  COMMIT;

SELECT SYSADM.JAVDRVN() FROM SYSIBM.SYSDUMMY1;

-- which returns "IBM Data Server Driver for JDBC and SQLJ"

SELECT SYSADM.JAVDRVV() FROM SYSIBM.SYSDUMMY1;

-- which returns 4.11.75 on the system used in this redbook.
```

If all the previous checks are OK, and you still have errors using the XSR routines, then you must look at the specific error codes and address them individually. Potential errors could include *SQLCODE -805 package not found* errors if you have not run the bind jobs, or Java errors if the Java environment is not defined correctly in the JAVASP.JSPENV file.

We did experience problems using XSR_COMPLETE that we resolved by adding a non-APF authorized data set to the STEPLIB concatenation in the WLM procedure used by XSR_COMPLETE.

2.2.4 z/OS XML system services

DB2 10 pureXML uses the z/OS XML system services for XML schema validation and XML parsing. These services are 100% eligible to be executed on a zIIP or zAAP processor. Figure 2-3 shows the z/OS XML system services processing flow. If the zAAP on zIIP feature is activated zAAP eligible z/OS XML system services workloads are eligible to be processed on a zIIP processor.

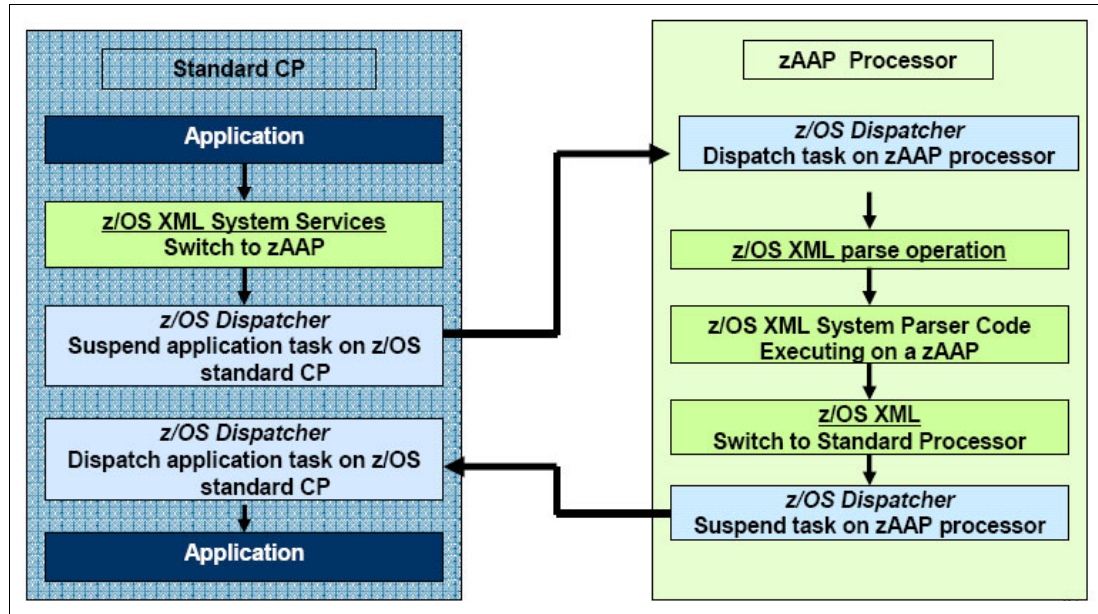


Figure 2-3 z/OS XML system services and zAAP processing flow

DB2 for z/OS pureXML XML parsing

DB2 pureXML invokes the z/OS XML system services for XML parsing. As a result, the XML parsing request becomes 100% zIIP or zAAP eligible, depending on whether the parsing or schema validation request is driven by DRDA® through a database access thread (DBAT) or through an allied DB2 thread.

Built in function DSN_XMLVALIDATE

In DB2 10, the SYSIBM.DSN_XMLVALIDATE function is provided inside the DB2 engine as a built-in function and uses z/OS XML System Services for XML validation. Thus, DSN_XMLVALIDATE invocations are 100% zIIP or zAAP eligible in DB2 10.

DB2 9 provided XML schema validation through the SYSFUN.DSN_XMLVALIDATE external UDF. The DB2 9 DSN_XMLVALIDATE UDF was executed in task control block (TCB) mode and did not use the z/OS XML system service for XML validation. Therefore, DSN_XMLVALIDATE invocations were neither zIIP nor zAAP eligible.

APARs PK90032 and PK90040 have enabled SYSIBM.DSN_XMLVALIDATE in DB2 9 and made that activity eligible for specialty engines..

For details on migrating to the new functions, see:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db2z10.doc.xml/db2z_udfdsxmlvalidatetobifdsxmlvalidate.htm

2.3 Choice of tools

Without naming individuals, the team that wrote this book included both persuasions of the user interface divide. At the half way point, the GUI² fan was amazed to find that the 3270 fan had deployed complex SQL/XML statements within SPUFI. The 3270 fan was bemused that this was considered strange.

We recognise that experienced DB2 professionals will have their own preferences for user interface, that will depend on the roles that they are trying to perform, the tasks that they are trying to execute, and their personal experiences. The good news is that there are plenty of tool choices on 3270 and on GUI.

2.3.1 3270 based tools

ISPF PDF, the DB2I panel, and SPUFI work just fine against DB2 tables with XML columns. If you are experienced in these tools and the ISPF editor, then you are able to perform all the DBA and development tasks that you need to.

You will frequently encounter the situation where your SQL and DDL statements contain an XML expression that is greater than 80 bytes long. In such circumstances, you can split an XPath expression over multiple lines and it will still run fine, as shown in Figure 2-4.

```

Menu Utilities Compilers Help
-----
BROWSE      XMLR3.XMLR3.OUT                      Line 00000000 Col 001 080
Command ==>                                     Scroll ==> PAGE
***** Top of Data *****
-----+-----+-----+-----+-----+-----+-----+-----+-----+
select xmlcast(xmlquery('declare default element namespace           00010003
"urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";                      00020003
$d/Document/BkToCstmrStmt                                           00030004
    /GrpHdr/MsgId'                                                  00040003
passing BK_TO_CSTMRT_STMT as "d") as varchar(35))                   00050003
from BK_TO_CSTMRT_STMT                                              00060003
-----+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+-----+-----+-----+
AAAASESS-FP-STAT001
...

```

Figure 2-4 Splitting an XPath expression over multiple lines in 3270 SPUFI session

DB2 Administration Tool for z/OS provides database administration facilities for DB2 objects containing XML columns. XML columns are just another data type that DB2 Administration Tool caters for.

2.3.2 GUI based tools

IBM offers a range of graphical based tools, mostly based on the eclipse framework, which support database administration and development activities with DB2 and pureXML.

Table 2-1 provides a summary of some of the tools which are most likely to be used with DB2 and pureXML.

² GUI stands for graphical user interface.

Table 2-1 IBM tools for DB2 administration and development with DB2 pureXML.

Tool	Overview of capabilities	Comments
Data Studio IDE	Data development (SQL, SQL/XML, stored procedures, Data Web Services etc...) XML: XML/Schema editor, XML mapper, Web services, Schema registration etc...	This is a download included in the DB2 licence for all platforms. It provides an eclipse-based environment for a wide range of DB2 development and administration activities, including XML.
Optim™ Development Studio	This is based on Data Studio, and includes extensions, such as <ul style="list-style-type: none"> ▶ pureQuery support ▶ XML validator ▶ XSD validator 	This is a chargeable tool. It includes additional facilities beyond Data Studio which are aimed more at a development user than a DBA.
InfoSphere™ Data Architect	Logical and physical data modeling - design databases, discover, relate, integrate and standardize diverse data assets. XML: Data Modeling, XML schema transformations	This is a chargeable tool. aimed at data modelers and architects. This tool provides logical and physical data modelling and schema development capabilities.
Rational® Developer for System z	Provides System z developers with tools for building traditional and composite applications in an SOA and Web 2.0 environment	This is a chargeable tool aimed at system z application developers. Contains overlap of XML related tools with the above products, and extended XML mapping and integration tools for system Z applications.

A comprehensive paper, "Tools and XML functionality for DB2 pureXML users" by Bryan Patterson, is available on IBM developerworks, which lists a wider range of IBM tools, and discusses the tasks that they support in the context of typical development and DBA roles. This paper can be accessed at the following URL:

<http://www.ibm.com/developerworks/data/library/techarticle/dm-1012xmltools/index.html>



Application scenario

In this chapter we introduce the application scenario which is based on requirement to process XML messages for business purposes. The message received contains bank statement information received from a financial institution. The scenario includes:

- ▶ Using an openly published XML standard as the basis for documents that we will be storing and manipulating in DB2.
- ▶ Integration with other systems via WebSphere® MQ to receive XML messages that are received by an organization's Enterprise Service Bus.
- ▶ Storing and processing the XML documents in DB2, using stored procedure, COBOL and JAVA programs which process XML and relational data together.
- ▶ Indexing and searching the XML tables, using standard SQL based query tools, to enable the business and audit requirements to be satisfied.

This chapter contains the following sections:

- ▶ Requirement for XML event logging and auditing
- ▶ Application scenario
- ▶ Application code samples

3.1 Requirement for XML event logging and auditing

Event logging, often used for auditing purposes, is one of the simplest use cases for DB2 pureXML.

Financial services companies are growing their usage of messaging and workflow systems at a phenomenal rate. Business processes are modelled and implemented, so that they may be executed as efficiently as possible. Enterprise Service Bus technology is used to link together all the systems that must be involved to complete these business processes. Messages are sent over an organization's Enterprise Service Bus during the automated execution of these business processes.

Whilst many different technologies and products may be used to implement these systems, there is often a common glue that binds them all together: XML.

An audit trail of some or all of the XML messages sent and received over the Enterprise Service Bus can be valuable for many reasons. For example

- ▶ An audit trail of messages describing high value financial transactions may provide a valuable source of information for monitoring and dashboard systems.
- ▶ Compliance requirements for certain applications may require that a very detailed audit trail be maintained.
- ▶ Technical problem resolution for messaging and workflow systems may be assisted by being able to trace the flow of messages through the execution of a business process.

Whatever the reasons for choosing to log event data in a persistent datastore, it is only worth logging the XML messages if you are subsequently able to read and analyze the XML messages with easy-to-use query and reporting tools.

DB2 pureXML provides an excellent platform for such event logging systems, because

- ▶ XML messages can be written to DB2 with minimal application development effort. No complex relational models need to be developed in order to log XML messages for subsequent auditing processing.
- ▶ The content of XML messages that have been stored in DB2 can be subsequently searched and accessed with ease using SQL/XML queries
- ▶ XML data elements used for searching and joining of data (such as userid, customerid, time and date) can be indexed so that queries can be optimized to access only those XML messages which are needed for the required task.

3.2 Application scenario

The application scenario (illustrated in Figure 3-1) in this book consists of the following aspects:

- ▶ We use the Bank To Customer Statement V2 (one of the ISO 20022 Universal financial industry message schemas) as the openly published XML standard for the XML documents that we save in DB2. ISO 20022 or UNIFI is the ISO Standard for Financial Services Messaging.
- ▶ We use the DB2 MQ Listener to receive XML messages from WebSphere MQ, and log them to DB2 using a native SQL stored procedure.
- ▶ We retrieve, manipulate and re-save those XML documents in DB2, using stored procedures, COBOL and JAVA programs. These programming examples show how the

XML data bindings work in these languages, and use the SQL/XML functions available within DB2.

- ▶ We create XML indexes for searching the XML audit tables, and show how normal SQL query tools can search these documents to enable the auditing requirements to be satisfied.

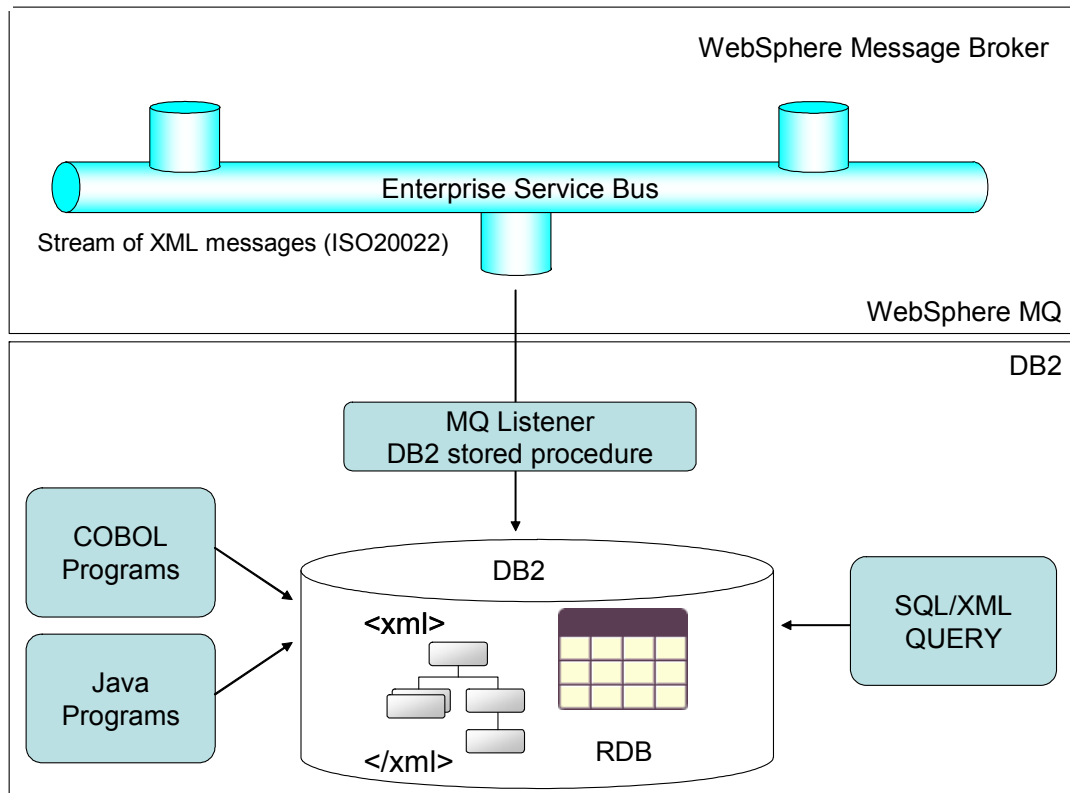


Figure 3-1 Application scenario

3.2.1 Using ISO 20022 with DB2 pureXML

A very common requirement when implementing XML data within DB2 is to work with an existing standard for XML messages. In our example, we have chosen to work with ISO 20022 (Universal financial industry message scheme).

The ISO 20022 standard provides the financial industry with a common set of messages in a standardized XML syntax.

- ▶ The current documentation for this standard is available at:
<http://www.iso20022.org/>
- ▶ The ISO 20022 document can be found at:
http://www.iso20022.org/catalogue_of_unifi_messages.page
- ▶ The description of the Bank To Customer Statement V2 can be downloaded from:
http://www.iso20022.org/documents/general/Payments_Maintenance_2009.zip
- ▶ The Schema of the Bank To Customer Statement V2 can be downloaded from:
<http://www.iso20022.org/documents/messages/camt/schemas/camt.053.001.02.zip>

The programming scenarios in this book are focussed on just one message type from the ISO20022 standard. (Bank To Customer Statement V2). We chose this message type because every reader will be familiar with the concept of a bank statement. For the ISO description of the message structure, you should download the description document referenced above. Briefly however, the XML message structure contains the following information.

- ▶ A group header, consisting of:
 - A unique message id
 - A message creation timestamp
 - Message pagination control information
- ▶ The statement itself, consisting of:
 - The statement id (duplicate of group header msgid element)
 - The statement creation timestamp (duplicate of group header element)
 - The period that the statement covers (from date and to date)
 - The account identification details
 - One or more account balances (multiple balance types exist)
 - All the transaction entries during the period of the statement

With DB2 pureXML it is easy to import the XML schemas that define an XML standard into the DB2 XML schema repository, so that all XML documents stored in DB2 can be automatically (or manually) validated against different versions of the XML schema standards.

The availability of an openly available set of industry standard XML schemas is a tremendously helpful productivity boost. Chapter 5, “Validating XML data” on page 73 provides worked examples of the ISO20022 standard schemas imported into DB2, so that the messages can be written into DB2 with full schema validation.

3.3 Application code samples

This book includes a number of programming samples that can be downloaded from the IBM Redbooks web site and are described in Appendix B, “Additional material” on page 273. The samples are provided in three groups:

- ▶ DB2 routines (stored procedures, UDFs etc...)
- ▶ COBOL programs
- ▶ Java programs

The DDL and schema samples are common to all the programming samples. However, each of the programming samples is independent of the other programming samples. So, if your interest is COBOL, you will be able to download and use the COBOL samples, without depending on Java or stored procedure samples.

The diagram in Figure 3-2 illustrates the flow between the various code samples presented in this book. The flow is explained in the following sections.

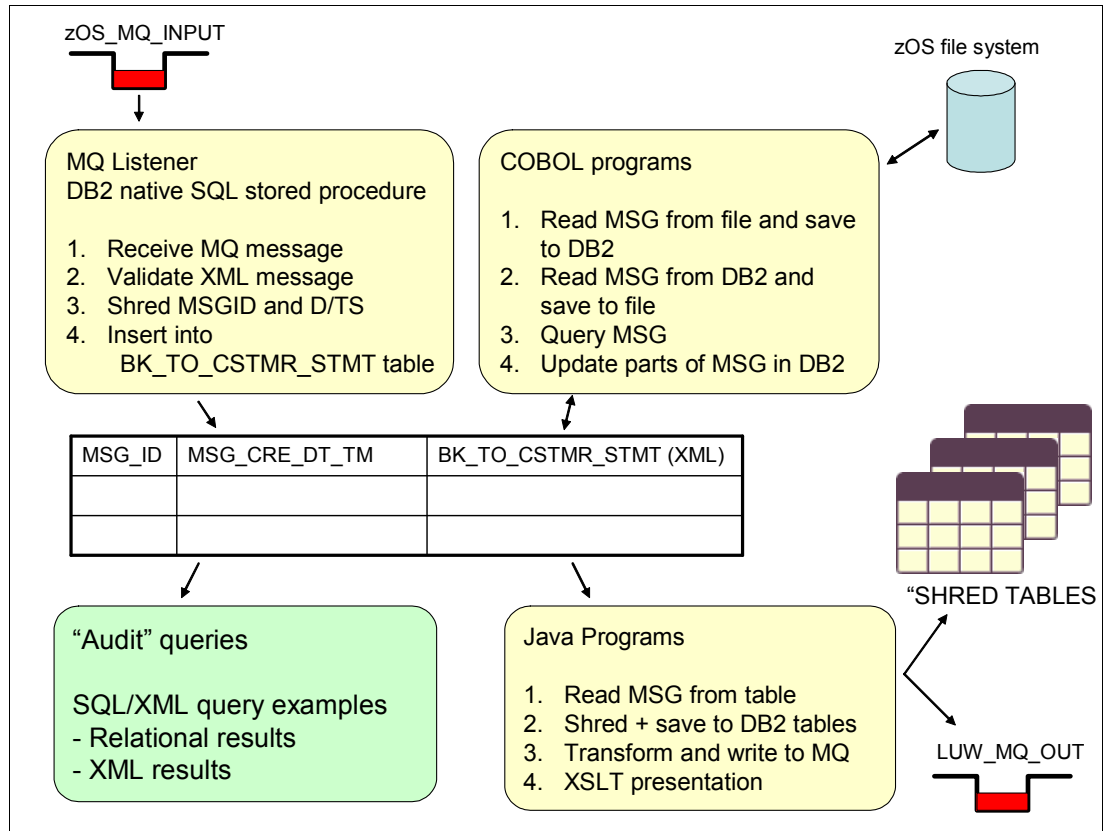


Figure 3-2 Four application code samples

3.3.1 DB2 SQL/XML programming pureXML

The amount of programming that happens within DB2 itself has been growing steadily as facilities like DB2 native SQL stored procedures have become more powerful. Stored procedures and user defined functions can be used in conjunction with external applications like COBOL and Java, where they provide re-usable routines that are productive to develop and perform exceptionally well within the DB2 engine.

A number of stored procedures are illustrated in Chapter 6, "DB2 SQL/XML programming" on page 87 covering the following tasks

- ▶ Using XML documents as input and output parameters to stored procedures which manipulate them (such as validating XML documents, shredding XML data elements into relational columns, transforming XML documents or sub-documents).
- ▶ Retrieving an XML messages from WebSphere MQ
- ▶ Using the DB2 MQ Listener to receive XML messages from WebSphere MQ, and automatically process them with a stored procedure
- ▶ Receiving XML messages from change data capture products, and using them to build XML documents that contain a history of data changes.

In addition to the stored procedures, this chapter also contains a number of SQL/XML examples that show how the SQL/XML language can be used to query the XML documents for auditing and other purposes.

The examples in this chapter are based on programs contained in a zip file called `storedproceduresamples.zip` as described in Appendix B, "Additional material" on page 273

and can be downloaded clicking on the additional material icon on the IBM Redbooks web site.

3.3.2 Using Java with DB2 pureXML

Chapter 7, “Using XML with Java” on page 129 covers Java programming in conjunction with DB2 pureXML. The examples in this chapter are based on programs are contained in a zip file called `javasamples.zip`, which can be downloaded from the additional materials page on the redbooks web site.

The functions of the Java programs are to

- ▶ Shred the previously saved XML message
- ▶ Query the message to produce a new XML document
- ▶ Output the new message to a WebSphere MQ queue
- ▶ Use the binary XML format to retrieve XML documents from DB2 to the IBM universal driver for JDBC and SQLJ.
- ▶ Perform XSLT transformations on the XML documents.

3.3.3 Using COBOL with DB2 pureXML

Chapter 8, “Using XML with COBOL” on page 155 covers COBOL programming in conjunction with DB2 pureXML. The examples in this chapter are based on programs contained in a zip file called `cobolsamples.zip`, which can be downloaded from the additional materials page on the redbooks web site.

A number of COBOL programs are presented in Chapter 8, “Using XML with COBOL” on page 155, they perform the following tasks:

- ▶ Read an XML message from a file and save it in DB2
- ▶ Select an XML message from DB2 and save it to a file
- ▶ Extract information from XML documents in DB2
- ▶ Update an XML document on a sub-document level

In addition we demonstrate what impact a schema change would have on the COBOL application and discuss how this compares to a relational implementation of the same database schema.

Finally, we look at some of the XML functionality available in native COBOL as a complement to the pureXML capabilities in DB2.



Creating and adding XML data

In this chapter we describe how to work with XML data.

We show how to create a table with an XML column and the storage structures that are used to handle the XML column. We also describe the multiversioning functions available with DB2 10.

We also shows how to retrieve information from the DB2 catalog tables for the base objects and the related XML objects, and demonstrates how an SQL INSERT statement can be used to move an XML document into an XML column.

We then include a brief overview of creating user indexes on an XML column.

This chapter contains the following sections:

- ▶ Creating and adding XML data
- ▶ Storage structure for XML data
- ▶ Multi-versioning concurrency control for XML
- ▶ Catalog queries to gather information
- ▶ Display database command
- ▶ Ingesting XML data
- ▶ XML indexes

4.1 Creation of tables with XML columns

To create tables with XML columns, you specify columns with the XML data type in the CREATE TABLE statement. A table can have one or more XML columns.

You do not specify a length when you define an XML column. There is no architectural limit on the size of an XML value in a database. However, textual XML data that is exchanged with a DB2 database is limited to 2 GB-1, so the effective limit of an XML column is 2 GB-1.

Like a LOB column, an XML column holds only a descriptor of the column. The data is stored separately.

Example 4-1 shows how you can define the BK_TO_CSTMR_STMT table referred to in 3.3, “Application code samples” on page 48 for the application scenario.

Example 4-1 BK_TO_CSTMR_STMT table with XML column

```
CREATE TABLE BK_TO_CSTMR_STMT
  (MSG_ID          VARCHAR(35),
   MSG_CRE_DT_TM  TIMESTAMP WITH TIMEZONE,
   BK_TO_CSTMR_STMT XML NOT NULL)
```

Since the IN clause is not specified, the table is created in an implicitly created partition-by-growth universal table space with a value of 256 for MAXPARTITIONS. The additional table space for the XML column is created using the default storage group SYSDEFLT in an implicitly created database.

4.2 Storage structure for XML data

The storage structure for XML data is similar to the storage structure for LOB data.

As with LOB data, the table that contains an XML column (the base table) is in a different table space from the table space which contains the XML data.

The storage structure depends on the type of table space that contains the base table.

We show the relationship between non-partitioned table space for base tables with XML columns and the corresponding XML table spaces and tables.

A base table space can be segmented as shown in Figure 4-1.

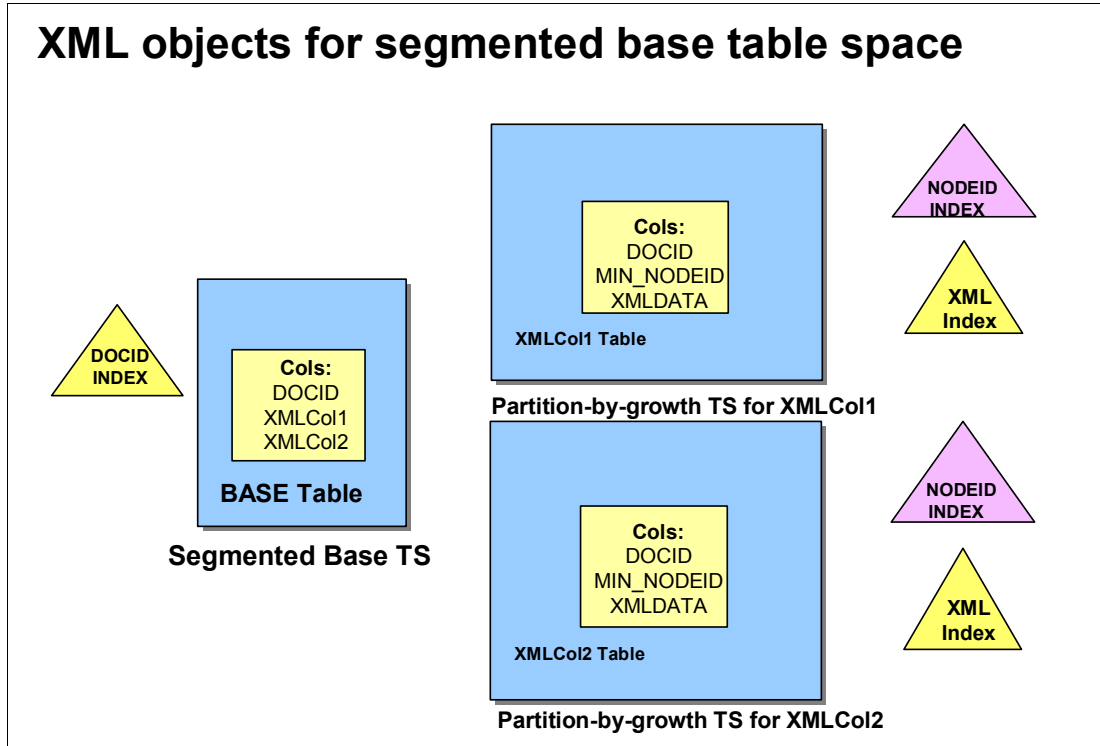


Figure 4-1 XML objects for segmented base table space

A base table space can be partition-by-growth shown in Figure 4-2.

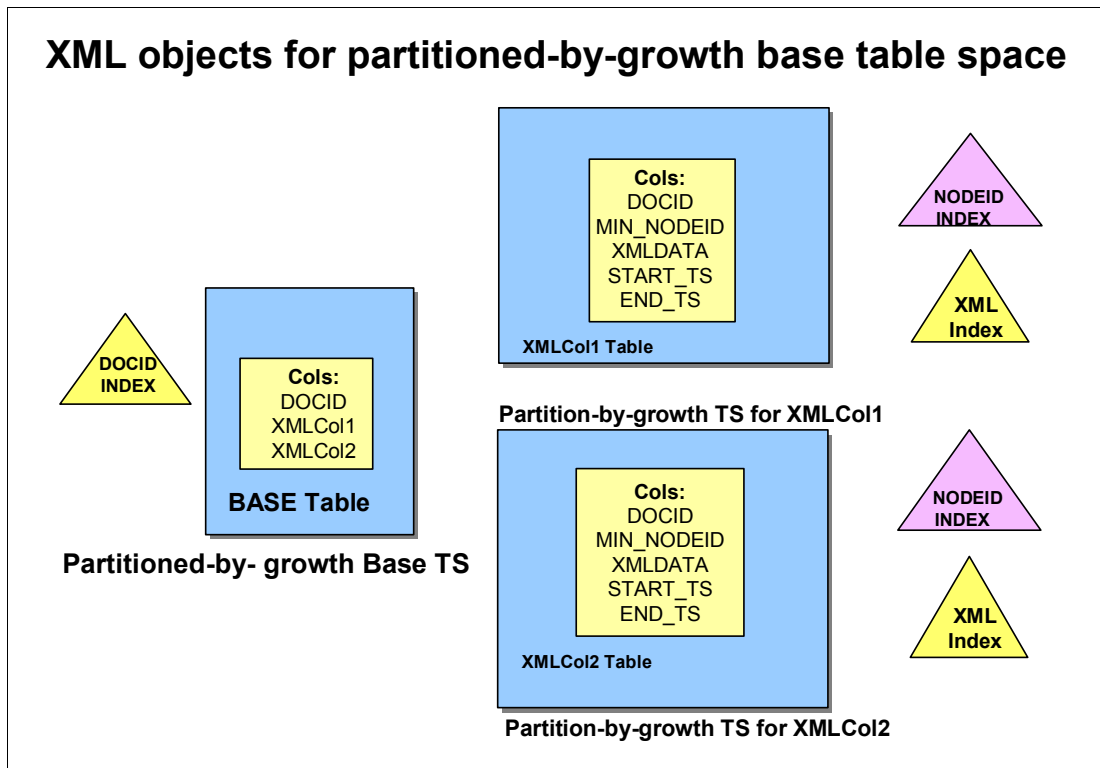


Figure 4-2 XML objects for partition-by-growth base table space

In both cases the XML table space is partition-by-growth. When you issue the CREATE TABLE statement which includes XML columns for a non-partitioned table or the ALTER TABLE statement to add an XML column to a non-partitioned table, the following objects are created implicitly by DB2 to support the XML columns:

- ▶ A column called DB2_GENERATED_DOCID_FOR_XML that is defined as NOT NULL. We refer to this column as DOCID column. DOCID uniquely represents each row. This column is hidden. For each table, DB2 only needs one DOCID column even if you would add additional columns with data type XML. This DOCID is defined as generated always, meaning that you cannot update the DOCID column.
- ▶ An XML indicator column in the base table for each XML column. The XML indicator column is treated like varying length column. An XML indicator column is VARCHAR(6) if the base table is segmented. It is stored in the base table in place of an XML column, and indicates whether the XML value for the column is null or zero length. The XML indicator column is VARCHAR(14) if the base table is in a partition-by growth universal table space. The extra 8 bytes are used to support multiple versions of XML documents.
- ▶ A unique index on the DOCID column. This index is known as a DOCID index. The DOCID index key is just the document ID itself pointing to the base table RID.
- ▶ An XML table space (partition-by-growth table space) per XML column which uses the Unicode UTF-8 encoding scheme
- ▶ An XML table with columns DOCID BIGINT, MIN_NODEID VARBINARY(128), and XMLDATA VARBINARY(15850) if the base table space is segmented. The XML table has two more columns START_TS and END_TS (used to support multiple versions of XML documents) if the base table space is partition-by-growth.
- ▶ An index on columns DOCID and XMLDATA in each XML table that DB2 uses to maintain document order, and map logical node ids to physical record IDs. This index is known as a NODEID index. The NODEID index is an extended, non partitioning index.
- ▶ DOCID and MIN_NODEID are used for row (XMLDATA) clustering and sort records in document order so prefetch will work.
- ▶ An XML document can span more than one partition. The base table space and the XML table space grow independently.

We show the relationship between partitioned table space for base tables with XML columns and the corresponding XML table spaces and tables.

A partitioned base table space can be classic partitioned as shown in Figure 4-3.

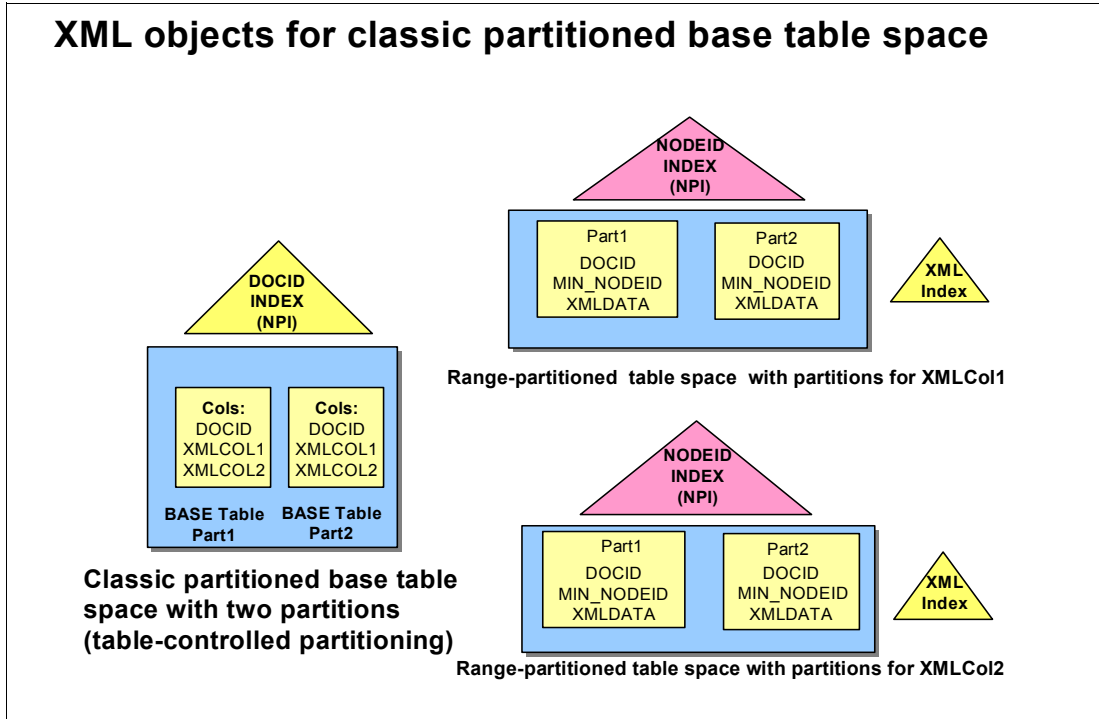


Figure 4-3 XML objects for classic partitioned base table space

A partitioned table space can be range partitioned shown in Figure 4-4.

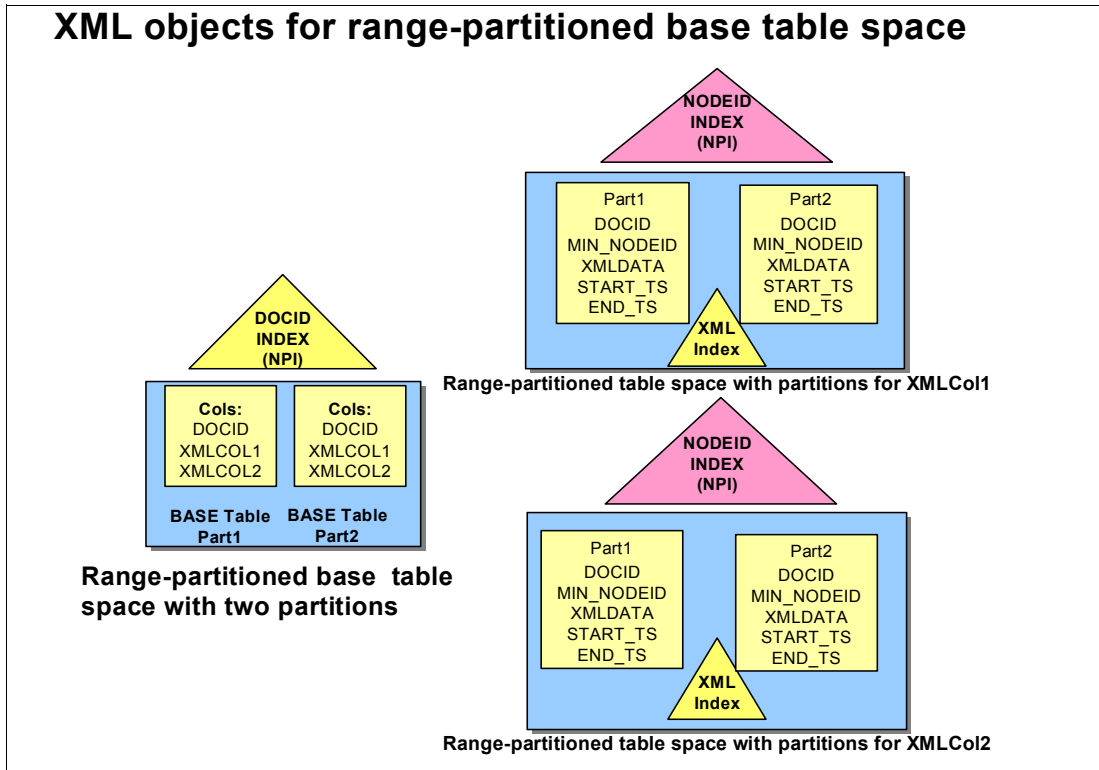


Figure 4-4 XML objects for range-partitioned base table space

When you issue the CREATE TABLE statement which includes XML columns for a partitioned table or the ALTER TABLE statement to add an XML column to a partitioned table, the following objects are created implicitly by DB2 to support the XML columns:

- ▶ A column called DB2_GENERATED_DOCID_FOR_XML that is defined as NOT NULL. We refer to this column as DOCID column. DOCID uniquely represents each row. This column is hidden. For each table, DB2 only needs one DOCID column even if you would add additional columns with data type XML. This DOCID is defined as generated always, meaning that you cannot update the DOCID column.
- ▶ A unique index on the DOCID column. This index is known as a DOCID index. The DOCID index key is just the document ID itself pointing to the base table RID. The DOCID index is a non-partitioning index.
- ▶ An XML indicator column in the base table for each XML column. The XML indicator column is treated like varying length column. An XML indicator column is VARCHAR(6). It is stored in the base table in place of an XML column, and indicates whether the XML value for the column is null or zero length. The XML indicator column is VARCHAR(14) if the base table is in a range partitioned universal table space. The extra 8 bytes are used to support multiple versions of XML documents.
- ▶ An XML table space (range-partitioned table space) per XML column which uses the Unicode UTF-8 encoding scheme
- ▶ Each XML table space contains one table with the corresponding number of parts. The XML table is partitioned only on the basis of the base row partition. Even though it is partitioned, the XML table space does not have limit keys. The XML data resides in the partition number that corresponds to the partition number of the base row. If a row changes partition in the base table, the XML document moves as well.
- ▶ An XML table with columns DOCID BIGINT, MIN_NODEID VARBINARY(128), and XMLDATA VARBINARY(15850) if the base table space is classic partitioned. The XML table has two more columns START_TS and END_TS (used to support multiple versions of XML documents) if the base table space is range-partitioned.
- ▶ An index on columns DOCID and XMLDATA in each XML table that DB2 uses to maintain document order, and map logical node ids to physical record IDs. This index is known as a NODEID index. The NODEID index is an extended, non-partitioning index.
- ▶ DOCID and MIN_NODEID are used for row (XMLDATA) clustering and sort records in document order so prefetch will work.

Important: The implicitly created XML table space is *always* a universal table space, either partition-by-growth or range-partitioned. If you want to have the ability to handle multiple versions of XML document make sure the base table space is a universal table space.

Figure 4-5 shows the storage structure of XML data when XML versions are not defined.

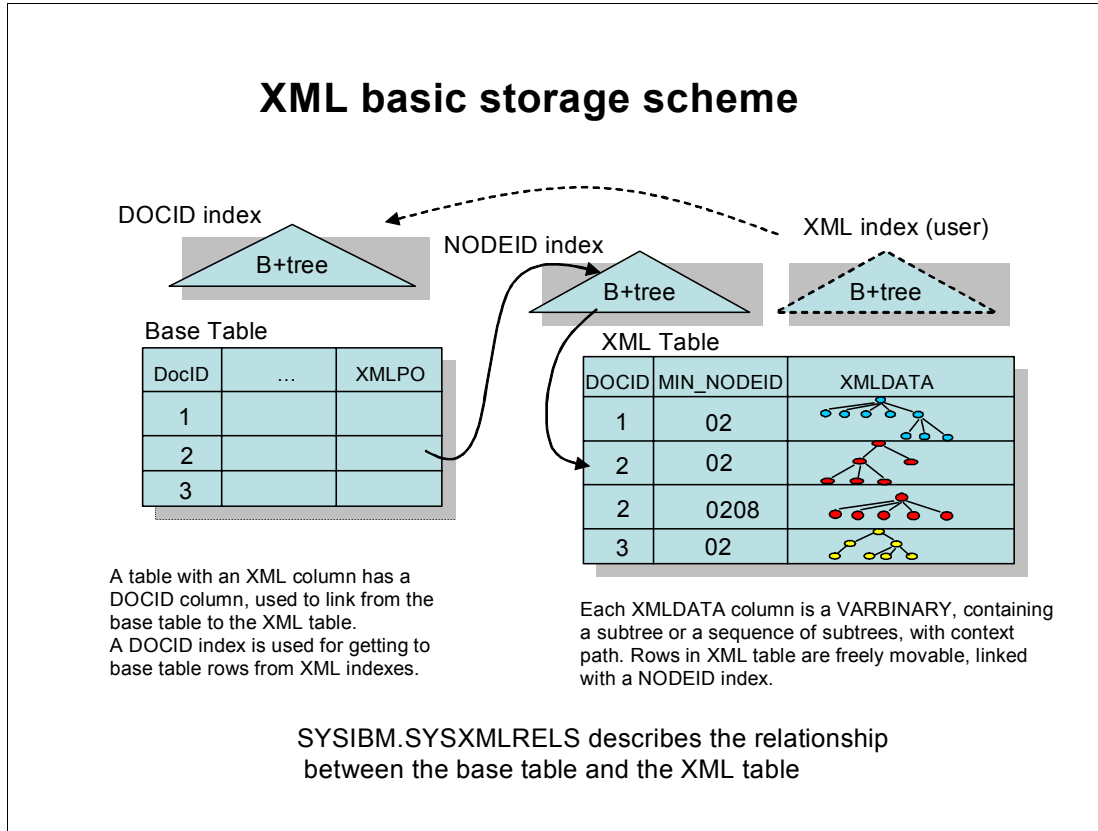


Figure 4-5 XML basic storage scheme

The XML column tells DB2 which NODEID index to search.

NODEID index is made up of the DOCID and a NODEID.

Key of DOCID+NODEID in the index tells DB2 which RID to get from XML tables.

NODEID Index may contain multiple entries per XML data row depending on the nodes grouped in the XML data row.

The NODEID index is created implicitly on the XML table for each XML column that is added to the base table. The XML node id uniquely identifies a node within a document.

To locate the corresponding XML column, DB2 uses the document ID from the base row and pairs that with an XML node id in the XML table, starting with the minimum node id, to search the NODEID index for the XML data.

The catalog table SYSIBM.SYSXMLRELS contains one row for each XML table that is created for an XML column to provide information about the relationship between XML column and XML table.

MIN_NODEID column

The node id for a node is the concatenation of local node ids contained in each node along the path from the root to the node. For each XML data record, there is a context node id that contains node ids from the root to the parent node for the nodes inside the record. The MIN_NODEID column of an XML table contains the minimum node id within the XMLDATA record in that row, which is concatenation of the context node id and the first node id in the

record body. DOCID and MIN_NODEID are used for clustering the rows belonging to a same document.

Note on XML indexes: XML indexes are user created indexes for achieving good performance.

4.3 Multi-versioning concurrency control for XML

DB2 supports multiple versions of an XML document in an XML column if the *base table space* for the table that contains the XML column is also a *universal table space* and created in DB2 10 NFM. All XML columns in the table support multiple versions.

Support note: If the base table space is not a universal table space, it does not support multiple XML versions. To convert the base table space from either segmented or partitioned to universal table space, you need to drop it and re-create it. ALTER and REORG are not sufficient in this case.

With XML versions, when you insert an XML document into an XML column, DB2 assigns a version number to the XML document. If the entire XML document is updated, DB2 creates a new version of the document in the XML table. If a portion of the XML document is updated, DB2 creates a new version of the updated portion. When DB2 uses XML versions, more data set space is required than when versions are not used. However, DB2 periodically deletes versions that are no longer needed. In addition, you can run the REORG utility against the XML table space that contains the XML document to remove unneeded versions. DB2 removes versions of a document when update operations that require the versions are committed, and when there are no readers that reference the unneeded versions.

XML versions note: XML versions are different from table space versions or index versions. The purpose of XML versions is to optimize concurrency and memory usage. The purpose of table space and index versions is to maximize data availability.

4.3.1 Example of improved concurrency with XML versions

The following example demonstrates how multiple XML versions can improve concurrency when the same XML documents are modified multiple times within the same transaction.

Suppose that table T1, which is in a universal table space, is defined like this:

```
CREATE TABLE T1(
  INT1 INT,
  XML1 XML,
  XML2 XML );
```

Table 4-1 shows the data in table T1.

Table 4-1 Data in table T1

INT1	XML1	XML2
350	<A1>111</A1>	<A2>aaa</A2>
100	<A1>111</A1>	<A2>aaa</A2>

INT1	XML1	XML2
250	<A1>111</A1>	<A2>aaa</A2>

An application performs SQL read operations that are represented by the following pseudocode:

```
EXEC SQL
  DECLARE CURSOR C1 FOR
  SELECT INT1, XML1
  FROM T1
  ORDER BY INT1
  FOR READ ONLY;
```

At the same time, another application performs SQL write operations that are represented by the following pseudocode:

```
EXEC SQL UPDATE T1
  SET XML1 = XMLPARSE(DOCUMENT '<B1>222</B1>');
EXEC SQL OPEN CURSOR C1; (Note: Cursor C1 is in another application as described)
EXEC SQL UPDATE T1
  SET XML1 = XMLPARSE(DOCUMENT '<C1>333</C1>');
EXEC SQL FETCH FROM C1 INTO :HVINT1, :HVXML1;
```

With multiple versions, the reading application does not need to hold a lock. Thus, the updating application can do its update operations without waiting for the reading application to finish. The reading application reads the old versions of the XML values, which are consistent data.

4.3.2 Example of improved storage usage with XML versions

The following example demonstrates how multiple XML versions can result in the use of less real storage when an XML document is the object of a self-referencing update operation.

Assume the same table T1 and data rows.

An application performs SQL operations that are represented by the following pseudocode:

```
EXEC SQL
  UPDATE T1
  SET XML1 = XML2, 1
  XML2 = XML1 2
  WHERE INT1 = 100;
EXEC SQL
  COMMIT 3;
```

The results of those operations are:

- 1.** When column XML1 is updated, DB2 stores the updated document as a new version in the XML table for column XML1. There are now two versions of the XML document for the second row of column XML1:
 - First version: <A1>111</A1>
 - Second version: <A2>aaa</A2>
- 2.** When column XML2 is updated, DB2 stores the updated document as a new version in the XML table for column XML2. There are now two versions of each XML document for the second row of column XML2:

- First version: <A2>aaa</A2>
- Second version: <A1>111</A1>

3. The update operations are committed. Thus, the old versions are no longer needed. DB2 deletes those versions from the XML tables for columns XML1 and XML2. (assuming no other readers are interested in reading these values).

Without multiple XML versions, DB2 needs to copy the original versions of the updated documents into memory, so that their values are not lost. For large XML documents, storage shortages might result.

4.3.3 Storage structure for XML data with versions

Figure 4-6 shows the storage structure of XML data to support XML versions.

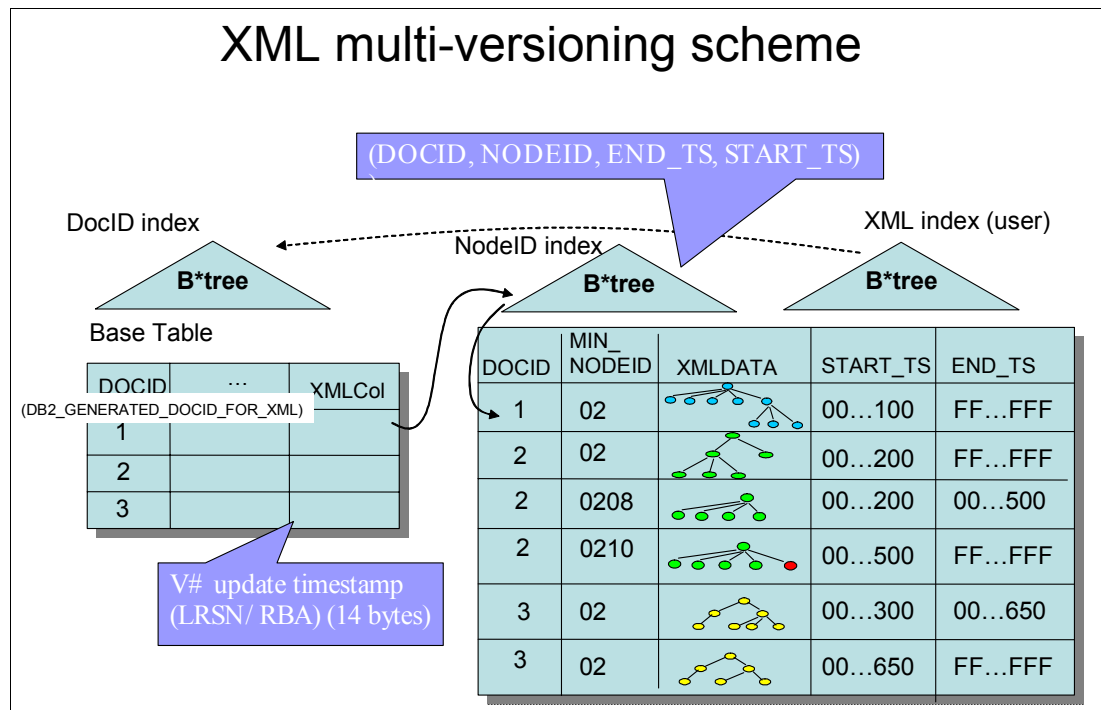


Figure 4-6 XML multi-versioning scheme

When you create a table with XML columns or ALTER a table to add XML columns, DB2 implicitly creates the following objects:

- ▶ A table space and table for each XML column. The data for an XML column is stored in the corresponding table.

DB2 creates the XML table space and table in the same database as the table that contains the XML column (the *base table*). The XML table space is in the Unicode UTF-8 encoding scheme.

If the base table contains XML columns that support XML versions, each XML table contains two more columns than an XML table for an XML column that does not support XML versions. Those columns are named START_TS and END_TS, and they have the BINARY(8) data type. START_TS contains the RBA or LRSN of the logical creation of an XML record. END_TS contains the RBA or LRSN of the logical deletion of an XML record. START_TS and END_TS identify the rows in the XML table that make up a version of an XML document.

Column START_TS represents the time when that row is created, and column END_TS represents the time when the row is deleted or updated. Column END_TS contains X'FFFFFFFFFFFFFFF' initially. To avoid compression causing update overflow, columns up to column END_TS are not compressed in the reordered row format.

- ▶ If the base table space supports XML versions, the length of the XML indicator column is eight bytes longer than the XML indicator column in a base table space that does not support XML versions. That is, 14 bytes instead of six bytes.
- ▶ A document ID column in the base table, named DB2_GENERATED_DOCID_FOR_XML, with data type BIGINT

We refer to this as DOCID column. The DOCID column holds a unique document identifier for the XML columns in a row. One DOCID column is used for all XML columns.

The DOCID column has the GENERATED ALWAYS attribute. Therefore, a value in this column cannot be NULL. However, it can be null for rows that existed before a table is altered to add an XML column.

- ▶ An index on the DOCID column.

This index is known as a document ID (or DOCID) index.

- ▶ An index on each XML table that DB2 uses to maintain document order, and map logical node ids to physical record IDs

This index is known as a NODEID index. The NODEID index is an extended, non-partitioning index.

If the base table space supports XML versions, the index key for the NODEID index contains two more columns than the index key for a node id index for a base table space that does not support XML versions. These are START_TS and END_TS.

Figure 4-7 gives a general idea of how DB2 handles multi-versioning for XML data.

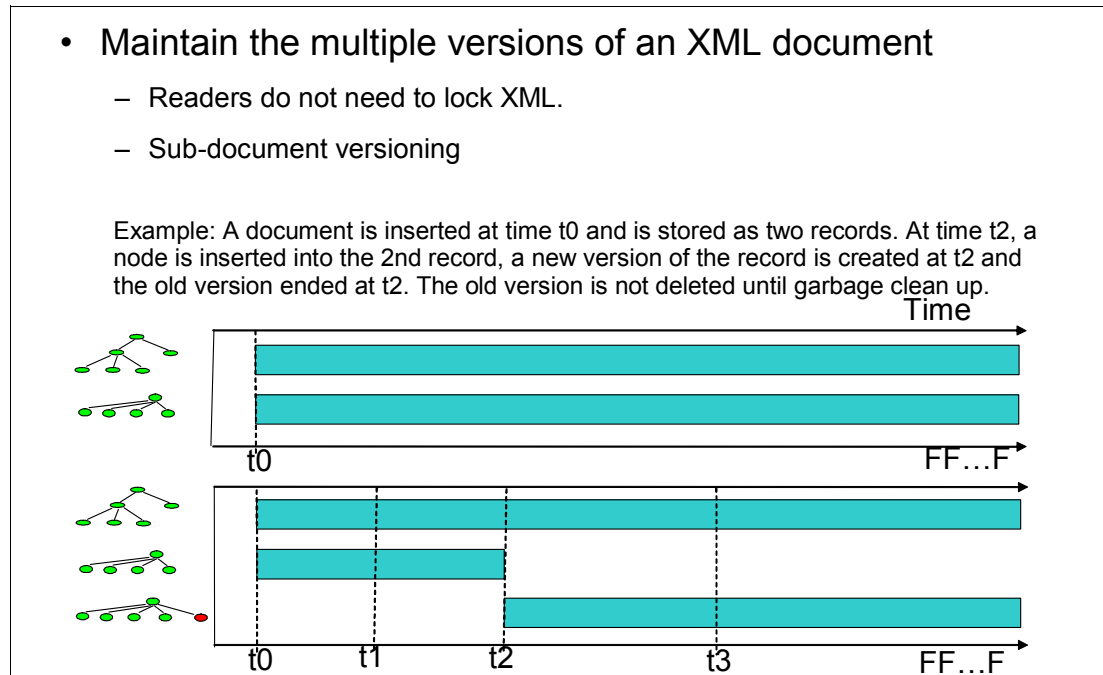


Figure 4-7 Multi-versioning for XML data

Each row in the XML auxiliary table is associated with two temporal attributes (start and end) to represent the period when the row exists. Start represents the time when that row is created, end represents the time when the row is deleted or expired.

For example, an XML document is stored as two rows in the XML auxiliary table at time t0, the second row is modified at t2, DB2 set that row expired at t2, create a new row representing the modified version with create time t2. The first row is not changed during this process.

A row in XML auxiliary table is never deleted until the garbage collector cleans it up.

When an XML document is deleted at time t2, all the records for that document are marked expired at t2. When a row of an XML document is updated, all the records for that document are marked expired at t2, the new document is inserted into XML auxiliary table with start time set to t2.

When a part of an XML document is modified, only the existing record (s) to be modified expire and a new version of those records is created.

Storage structure note: This storage structure is possible only in new-function mode and for universal table spaces. Storage structure for multi-versioning is a prerequisite for several other XML enhancements such as:

- ▶ Access of currently committed data
- ▶ “As Of” for time oriented data
- ▶ XML Update with XMLMODIFY
- ▶ Removing restrictions for SELECT FROM UPDATE//DELETE for XML

Figure 4-8 shows the XML locking scheme in DB2 10.

XML locking scheme with multi-versioning			
SQL	Base Page/Row Lock (Business as usual)	XML Lock	XML Table space Page Lock
INSERT	x page/row lock	x lock, release at commit	Page latch (and optional P-Lock)
UPDATE/DELETE	u->x, s->x, x stays x	x lock, release at commit	Page latch (and optional P-Lock)
SELECT UR, SELECT CS- CURRENT DATA NO	None	Conditional lock	
SELECT CS- CURRENT DATA YES no workfile	s page/row lock, release on next row fetch		
SELECT CS- CURRENT DATA YES workfile	s page/row lock, release on next row fetch		
SELECT UR, SELECT CS- CURRENT DATA NO, SELECT CS- CURRENT DATA YES with Multirow fetch and dynamic scrolling	s page/row lock on rowset, release on next fetch		
SELECT RR, SELECT RS	s page/row lock		

Figure 4-8 XML Locking scheme with multi-versioning

4.4 Catalog queries to gather information

You can get hold of the implicitly created database and table space names for the table BK_TO_CSTMR_STMT from the catalog table SYSIBM.SYSTABLES as shown in **Query 1** in Example 4-2.

Once you know the name of the implicitly created database, you can retrieve information about the default storage group, default buffer pool used for data, and default buffer pool used for indexes from the catalog table SYSIBM.SYSDATABASE as shown in **Query 2** in Example 4-2. The value 'Y' in column IMPLICIT indicates the database is implicitly created.

Once you know the name of the implicitly created database, you can retrieve information about table spaces in this database from the catalog table SYSIBM.SYSTABLESPACE as shown in **Query 3** in Example 4-2.

There are two rows in this table:

- ▶ The first row is for the base table space BKRTORCS. The value 'G' in column TYPE indicates this is a partition by growth universal table space, currently has one partition (indicated by column PARTITIONS) and can grow to a maximum of 256 partitions (indicated by column MAXPARTITIONS). The value 'Y' in column IMPLICIT indicates the table is implicitly created. SEGSIZE used is 32. The maximum data set size is 4 GB, indicated by column DSSIZE in kilobytes.
- ▶ The second row is for the XML table space XBKR0000. The value 'P' in column TYPE indicates this is an implicit table space created for XML columns, currently has one partition (indicated by column PARTITIONS) and can grow to a maximum of 256 partitions (indicated by column MAXPARTITIONS). The value 'Y' in column IMPLICIT indicates the table is implicitly created. SEGSIZE used is 32. The maximum data set size is 4 GB, indicated by column DSSIZE in kilobytes.

Note: The name of the XML table space always starts with X and it is always a universal table space. In this case it is partition-by-growth because the base table space is partition-by-growth, and is created in the same database as the base table space.

Example 4-2 Catalog Queries (1 of 3)

-- Query 1

```
SELECT SUBSTR(NAME,1,20) AS NAME,
       SUBSTR(DBNAME,1,10) AS DBNAME,
       SUBSTR(TSNAME,1,10) AS TSNAME
FROM SYSIBM.SYSTABLES
WHERE NAME='BK_TO_CSTMR_STMT' AND CREATOR=USER;
-----+-----+-----+-----+-----+-----+-----+-----+
NAME                DBNAME            TSNAME
-----+-----+-----+-----+-----+-----+-----+-----+
BK_TO_CSTMR_STMT    DSN00242          BKRTORCS
DSNE610I NUMBER OF ROWS DISPLAYED IS 1
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100
-----+-----+-----+-----+-----+-----+-----+-----+
```

-- Query 2

```
SELECT SUBSTR(NAME,1,10) AS NAME,
       SUBSTR(CREATOR,1,10) AS CREATOR,
```

```

SUBSTR(STGROUP,1,10) AS STGROUP,
BPOOL,INDEXBP,IMPLICIT
FROM SYSIBM.SYSDATABASE
WHERE NAME='DSN00242' ;
-----+-----+-----+-----+-----+-----+-----
NAME          CREATOR    STGROUP    BPOOL      INDEXBP    IMPLICIT
-----+-----+-----+-----+-----+-----+-----
DSN00242     SYSIBM     SYSDEFLT   BPO         BPO         Y
DSNE610I NUMBER OF ROWS DISPLAYED IS 1
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100
-----+-----+-----+-----+-----+-----+-----

```

-- Query 3

```

SELECT SUBSTR(NAME,1,10) AS NAME,
SUBSTR(CREATOR,1,10) AS CREATOR,
SUBSTR(DBNAME,1,10) AS DBNAME,
STATUS,TYPE,SEGSIZE,PARTITIONS,
MAXPARTITIONS AS MAXP,
DSSIZE,IMPLICIT
FROM SYSIBM.SYSTABLESPACE
WHERE DBNAME = 'DSN00242'
AND CREATOR = USER ;
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
NAME          CREATOR  DBNAME    STATUS  TYPE  SEGSIZE  PARTITIONS  MAXP  DSSIZE  IMPLICIT
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
BKRTORCS     XMLR4    DSN00242  A      G      32        1    256  4194304  Y
XBKR0000     XMLR4    DSN00242  A      P      32        1    256  4194304  Y
DSNE610I NUMBER OF ROWS DISPLAYED IS 2
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----

```

Once you know the name of the implicitly created database, you can retrieve information about tables in this database from the catalog table SYSIBM.SYSTABLES as shown in **Query 4** in Example 4-3.

There are two rows in this table:

- ▶ The first row is for the base table BK_TO_CSTMTR_STMT (the value 'T' in column TYPE indicates this). The blank value in column STATUS indicates that the table has no unique constraint (primary key or unique key) and the definition of the table is complete. The blank value in column TABLESTATUS indicates that the table definition is complete.
- ▶ The second row is for the XML table XBK_TO_CSTMTR_STMT. The value 'P' in column TYPE indicates this is an Implicit table created for XML columns. The blank value in column STATUS indicates that the table has no unique constraint (primary key or unique key) and the definition of the table is complete. The blank value in column TABLESTATUS indicates that the table definition is complete.

The name of the XML table always starts with X.

You can get information about the columns in both the base table and XML table from SYSIBM.SYSCOLUMNS as shown in **Query 5** in Example 4-3. The first three rows are for the three columns defined in the base table. The fourth row is the DOCID column generated by DB2. The next three rows are for columns defined in the XML table. The last two rows are for the two extra columns defined in the XML table to support multiple versions for XML documents and are present because the base table is in partition-by-growth universal table space.

*Example 4-3 Catalog Queries (2 of 3)***-- Query 4**

```

SELECT SUBSTR(NAME,1,20) AS NAME,
       SUBSTR(CREATOR,1,10) AS CREATOR,
       TYPE,
       SUBSTR(DBNAME,1,10) AS DBNAME,
       SUBSTR(TSNAME,1,10) AS TSNAME,
       STATUS, TABLESTATUS
FROM SYSIBM.SYSTABLES
WHERE DBNAME = 'DSN00242'
      AND CREATOR = USER ;

```

```

-----+-----+-----+-----+-----+-----+-----+
NAME          CREATOR    TYPE  DBNAME    TSNAME    STATUS TABLESTATUS
-----+-----+-----+-----+-----+-----+-----+
BK_TO_CSTMR_STMT  XMLR4      T    DSN00242  BKRTORCS
XBK_TO_CSTMR_STMT XMLR4      P    DSN00242  XBKR0000
DSNE610I NUMBER OF ROWS DISPLAYED IS 2
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100
-----+-----+-----+-----+-----+-----+-----+

```

-- Query 5

```

SELECT SUBSTR(NAME,1,30) AS NAME,
       SUBSTR(TBNAME,1,20) AS TBNAME,
       SUBSTR(TBCREATOR,1,10) AS TBCREATOR,
       SUBSTR(COLTYPE,1,8) AS COLTYPE,
       LENGTH, COLNO
FROM SYSIBM.SYSCOLUMNS
WHERE TBCREATOR = USER
      AND TBNAME IN ('BK_TO_CSTMR_STMT', 'XBK_TO_CSTMR_STMT')
ORDER BY TBNAME, COLNO;

```

```

-----+-----+-----+-----+-----+-----+-----+
NAME          TBNAME          TBCREATOR  COLTYPE  LENGTH  COLNO
-----+-----+-----+-----+-----+-----+-----+
MSG_ID        BK_TO_CSTMR_STMT XMLR4      VARCHAR  35      1
MSG_CRE_DT_TM BK_TO_CSTMR_STMT XMLR4      TIMESTMP 12      2
BK_TO_CSTMR_STMT BK_TO_CSTMR_STMT XMLR4      XML      14      3
DB2_GENERATED_DOCID_FOR_XML BK_TO_CSTMR_STMT XMLR4      BIGINT   8       4
DOCID         XBK_TO_CSTMR_STMT XMLR4      BIGINT   8       1
MIN_NODEID   XBK_TO_CSTMR_STMT XMLR4      VARBIN   128     2
XMLDATA      XBK_TO_CSTMR_STMT XMLR4      VARBIN  15850   3
START_TS     XBK_TO_CSTMR_STMT XMLR4      BINARY   8       4
END_TS       XBK_TO_CSTMR_STMT XMLR4      BINARY   8       5
DSNE610I NUMBER OF ROWS DISPLAYED IS 4
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100
-----+-----+-----+-----+-----+-----+-----+

```

The relationship between the base table and XML table is described in the catalog table SYSIBM.SYSXMLRELS and this information can be retrieved as shown in **Query 6** in Example 4-4. Column XMLRELOBID shows internal identifier of the relationship between the base table and the XML table.

Information about the DOCID and NODEID indexes can be retrieved from the catalog table SYSIBM.SYSINDEXES as shown in **Query 7** in Example 4-4.

Column IX_EXTENSION_TYPE shows 'N' in the first row indicating this row is for the NODEID index which is an extended index. Column IX_EXTENSION_TYPE has blanks in the second row indicating this row is for the DOCID index which is a simple index.

Example 4-4 Catalog Queries (3 of 3)

-- Query 6

```
SELECT SUBSTR(TBOWNER,1,10) AS TBOWNER,
       SUBSTR(TBNAME,1,20) AS TBNAME,
       SUBSTR(COLNAME,1,20) AS COLNAME,
       SUBSTR(XMLTBOWNER,1,10) AS XMLTBOWNER,
       SUBSTR(XMLTBNAME,1,20) AS XMLTBNAME,
       XMLRELOBID
FROM SYSIBM.SYSXMLRELS
WHERE TBOWNER = USER
     AND TBNAME = 'BK_TO_CSTMRTMT';
```

TBOWNER	TBNAME	COLNAME	XMLTBOWNER	XMLTBNAME	XMLRELOBID
XMLR4	BK_TO_CSTMRTMT	BK_TO_CSTMRTMT	XMLR4	XBK_TO_CSTMRTMT	7

DSNE610I NUMBER OF ROWS DISPLAYED IS 1
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100

--Query 7

```
SELECT SUBSTR(NAME,1,30) AS NAME,
       SUBSTR(CREATOR,1,10) AS CREATOR,
       SUBSTR(TBNAME,1,20) AS TBNAME,
       SUBSTR(TBCREATOR,1,10) AS TBCREATOR,
       SUBSTR(DBNAME,1,10) AS DBNAME,
       SUBSTR(INDEXSPACE,1,20) AS INDEXSPACE,
       IX_EXTENSION_TYPE AS IXET
FROM SYSIBM.SYSINDEXES
WHERE DBNAME = 'DSN00242'
     AND CREATOR = USER ;
```

NAME	CREATOR	TBNAME	TBCREATOR	DBNAME	INDEXSPACE	IXET
I_NODEIDXBK_TO_CSTMRTMT	XMLR4	XBK_TO_CSTMRTMT	XMLR4	DSN00242	IRNODEID	N
I_DOCIDBK_TO_CSTMRTMT	XMLR4	BK_TO_CSTMRTMT	XMLR4	DSN00242	IRDOCIDB	

DSNE610I NUMBER OF ROWS DISPLAYED IS 2

4.5 Display database command

If you issue the DB2 -DISPLAY DATABASE command, you can see all the base objects and the XML objects associated with the database as shown in Figure 4-9 on page 67.

TYPE is TS for a table space, IX for an index space, and XS for a XML table space.

PART is the partition number. Since we have one partition for the partition-by-growth table space for both the base table and the XML table, PART shows 00001. Over time the table space can grow to more partitions and that is why there are two lines for each of the base and XML table spaces, at present with no value.

For non-partitioned indexes, it is the logical partition number preceded by the character L (for example, L0001). This is the case for the DOCID and NODEID indexes.

```

-DISPLAY DB(DSN00242)

DSNT360I  -DBOB *****
DSNT361I  -DBOB *   DISPLAY DATABASE SUMMARY
           *   GLOBAL
DSNT360I  -DBOB *****
DSNT362I  -DBOB      DATABASE = DSN00242  STATUS = RW
           DBD LENGTH = 4028

DSNT397I  -DBOB
NAME      TYPE PART  STATUS          PHYERRLO PHYERRHI CATALOG  PIECE
-----
BKRTORCS TS      0001 RW
BKRTORCS TS              RW
XBKR0000 XS      0001 RW
XBKR0000 XS              RW
IRDOCIDB IX     L0001 RW
IRDOCIDB IX       L*   RW
IRNODEID IX     L0001 RW
IRNODEID IX       L*   RW
***** DISPLAY OF DATABASE DSN00242 ENDED *****
DSN9022I  -DBOB DSNTDDIS 'DISPLAY DATABASE' NORMAL COMPLETION
***

```

Figure 4-9 -DISPLAY DATABASE command output

4.6 Ingesting XML data

You can use different techniques to ingest XML data. There are different forms of the INSERT statement and the LOAD utility.

In this section, we show an example of the SQL INSERT statement using a string literal to insert rows into a table that contains XML columns. This form of INSERT is suitable for small documents. The host variable or file reference versions of INSERT are applicable for any length.

Example 4-5 shows the successful insertion of the shorter version of the Message Received XML document for our application scenario shown in Example A-2 on page 269.

Example 4-5 Using the SQL INSERT statement to insert XML document to an XML column

```

INSERT INTO XMLR4.BK_TO_CSTMR_STMT(BK_TO_CSTMR_STMT)
values('<?xml version="1.0" encoding="UTF-8" ?>
  <Document xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02">
  <BkToCstmrStmt>
    <GrpHdr>
      <MsgId>AAAASESS-FP-STAT001</MsgId>
      <CreDtTm>2010-10-18T17:00:00+01:00</CreDtTm>
    </GrpHdr>
    <Stmt>
  </Id>AAAASESS-FP-STAT001</Id>

```

```

<CreDtTm>2010-10-18T17:00:00+01:00</CreDtTm>
  <FrToDt>
    <FrDtTm>2010-10-18T08:00:00+01:00</FrDtTm>
    <ToDtTm>2010-10-18T17:00:00+01:00</ToDtTm>
  </FrToDt>
<Acct>
  <Id>
    <Othr>
      <Id>50000000054910000003</Id>
    </Othr>
  </Id>
  <Ownr>
    <Nm>FINPETROL</Nm>
  </Ownr>
  <Svcr>
    <FinInstnId>
      <Nm>AAAA BANKEN</Nm>
    </FinInstnId>
  </Svcr>
</Acct>
<Bal>
  <Tp>
<CdOrPrtry>
  <Cd>OPBD</Cd>
  </CdOrPrtry>
  </Tp>
  <Amt Ccy="SEK">500000</Amt>
  <CdtDbtInd>CRDT</CdtDbtInd>
</Bal>
<Bal>
  <Tp>
    <CdOrPrtry>
      <Cd>CLBD</Cd>
    </CdOrPrtry>
  </Tp>
  <Amt Ccy="SEK">435678.50</Amt>
  <CdtDbtInd>CRDT</CdtDbtInd>
</Bal>
<Ntry>
<Amt Ccy="SEK">105678.50</Amt>
  <BookgDt>
    <DtTm>2010-10-18T13:15:00+01:00</DtTm>
  </BookgDt>
  <AcctSvcrRef>AAAASESS-FP-CN_98765/01</AcctSvcrRef>
</Ntry>
<Ntry>
  <Amt Ccy="SEK">200000</Amt>
  <BookgDt>
    <DtTm>2010-10-18T10:15:00+01:00</DtTm>
  </BookgDt>
  <AcctSvcrRef>AAAASESS-FP-ACCR-01</AcctSvcrRef>
</Ntry>
<Ntry>
  <Amt Ccy="SEK">30000</Amt>
  <BookgDt>

```

```

        <DtTm>2010-10-18T15:15:00+01:00</DtTm>
    </BookgDt>
    <AcctSvcrRef>AAAASESS-FP-CONF-FX</AcctSvcrRef>
  </Ntry>
</Stmt>
<Stmt>
  <Id>AAAASESS-FP-STAT002</Id>
  <CreDtTm>2010-10-17T17:00:00+01:00</CreDtTm>
  <FrToDt>
    <FrDtTm>2010-10-17T08:00:00+01:00</FrDtTm>
    <ToDtTm>2010-10-17T17:00:00+01:00</ToDtTm>
  </FrToDt>
  <Acct>
    <Id>
      <Othr>
        <Id>50000000054910000004</Id>
      </Othr>
    </Id>
    <Ownr>
      <Nm>FINPETROL</Nm>
    </Ownr>
    <Svcr>
      <FinInstnId>
        <Nm>AAAB BANKEN</Nm>
      </FinInstnId>
    </Svcr>
  </Acct>
  <Bal>
    <Tp>
      <CdOrPrtry>
        <Cd>OPAV</Cd>
      </CdOrPrtry>
    </Tp>
    <Amt Ccy="SEK">500300</Amt>
    <CdtDbtInd>CRDT</CdtDbtInd>
  </Bal>
  <Bal>
    <Tp>
      <CdOrPrtry>
        <Cd>FWAV</Cd>
      </CdOrPrtry>
    </Tp>
    <Amt Ccy="SEK">435478.50</Amt>
    <CdtDbtInd>CRDT</CdtDbtInd>
  </Bal>
  <Ntry>
    <Amt Ccy="SEK">105378.50</Amt>
    <CdtDbtInd>CRDT</CdtDbtInd>
    <ValDt>
      <Dt>2010-10-17</Dt>
    </ValDt>
    <AcctSvcrRef>AAAASESS-FP-CN_98764/01</AcctSvcr
  </Ntry>
</Ntry>
<Amt Ccy="SEK">200100</Amt>

```

```

    <ValDt>
      <Dt>2010-10-17</Dt>
    </ValDt>
    <AcctSvcrRef>AAAASESS-FP-ACCR-02</AcctSvcrRef>
  </Ntry>
  <Ntry>
    <Amt Ccy="SEK">30020</Amt>
    <BookgDt>
      <DtTm>2010-10-17T15:15:00+01:00</DtTm>
    </BookgDt>
    <AcctSvcrRef>AAAASESS-FP-CONF-FY</AcctSvcrRef>
  </Ntry>
</Stmt>
</BkToCstmrStmt>
</Document>');

```

```

DSNE615I NUMBER OF ROWS AFFECTED IS 1
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 0

```

This technique is quite cumbersome for large XML documents. We show how to use the LOAD utility in such cases. LOAD utility is discussed in Chapter 9, “Utilities with XML” on page 181.

4.7 XML indexes

An XML index can be used to improve the efficiency of queries on XML documents that are stored in an XML column.

Just like you define relational indexes on selected columns of a relational table, you define XML indexes on selected elements and attributes within a single XML column of a table. In particular, XML indexes in DB2 do not automatically index all the values in an XML column, but only the ones that you choose. Although you can choose to index all elements and attributes, you should typically index just those elements and attributes that are frequently used in predicates and join conditions.

Instead of providing access to the beginning of a document, index entries in an XML index provide access to nodes within the document by creating index keys based on XML pattern expressions. Because multiple parts of a XML document can satisfy an XML pattern, DB2 might generate multiple index keys when it inserts values for a single document into the index.

You create an XML index using the CREATE INDEX statement, and drop an XML index using the DROP INDEX statement. The GENERATE KEY USING XMLPATTERN clause you include with the CREATE INDEX statement specifies what you want to index.

Some of the keywords used with the CREATE INDEX statement for indexes on non-XML columns do not apply to indexes over XML data.

Example 4-6 shows an example of creating an index on the sample XML document for the application scenario.

Example 4-6 XML index on DtTm elements

```

CREATE INDEX IXMLNTRY
ON BK_TO_CSTMR_STMT(BK_TO_CSTMR_STMT) 1
GENERATE KEY USING XMLPATTERN 2

```

```
'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
/Document/BkToCstmrStmnt/Stmnt/Ntry/BookgDt/DtTm'
AS SQL TIMESTAMP 3
```

Note the following in Example 4-6:

- 1.** The XML index is defined on the BK_TO_CSTMTR_STMT column of the BK_TO_CSTMTR_STMT table. BK_TO_CSTMTR_STMT must be of the XML data type.
- 2.** The GENERATE KEY USING XMLPATTERN clause provides information about what you want to index. This clause is called an XML index specification. The XML index specification contains an XML pattern clause. The XML pattern clause in this example indicates that you want to index the values of the DtTm element. The namespace declaration is necessary because the XML documents have a namespace declaration.
- 3.** AS SQL TIMESTAMP indicates that indexed values are stored as TIMESTAMP values.

Only one index specification clause is allowed in a CREATE INDEX statement. However, you can create multiple XML indexes on an XML column.

Every XML pattern expression that you specify in a CREATE INDEX statement must be associated with a data type. The data type must be VARCHAR, DECFLOAT, DATE, or TIMESTAMP.

You can interpret the result of pattern expression as multiple data types. For example, the value 123 has a character representation, but it can also be interpreted as the number 123. You can create different indexes on the same pattern expression with different data types, so that the data can be indexed, regardless of its data type.

If you validate your XML documents against an XML schema, ensure that the data type specifications in the XML schema match the data types that you use for your indexes.

The UNIQUE keyword in XML index definitions has a similar but slightly different meaning than it does for relational index definitions.

- ▶ For relational indexes, the UNIQUE keyword in the CREATE INDEX statement enforces uniqueness across all rows in a table.
- ▶ For indexes over XML data, the UNIQUE keyword enforces uniqueness across all documents in an XML column.

For an XML index, DB2 enforces uniqueness for:

- ▶ The data type of the index.
- ▶ The XML path to a node.
- ▶ The value of the node after the XML value has been cast to the SQL data type that is specified for the index.

Because rounding can occur during conversion of an index key value to the specified data type for the index, multiple values that appear to be unique in the XML document might rarely result in duplicate key errors.



Validating XML data

In this chapter we discuss details of user-controlled and automatic validation of XML documents. The sections are:

- ▶ XML schema validation
- ▶ XML type modifier
- ▶ Automatic validation
- ▶ User-controlled validation
- ▶ Determining whether an XML document has been validated

5.1 XML schema validation

XML schema validation is the process of determining whether the structure, content, and data types of an XML document are valid according to an XML schema.

In addition, XML schema validation strips ignorable whitespace from the input document.

There are two ways that you can validate an XML document in DB2 10:

- ▶ Automatically, by including an XML type modifier in the XML column definition in a CREATE TABLE or ALTER TABLE statement. When a column has an XML type modifier, DB2 implicitly validates documents that are inserted into the column or when documents in the column are updated.
- ▶ User-controlled, by executing the DSN_XMLVALIDATE built-in function when you insert a document into an XML column or update a document in an XML column or before selecting back (not necessarily into a table).

Validation is optional when you insert data into an XML column with no XML type modifier. Validation is mandatory when you insert data into an XML column with an XML type modifier.

We examine both methods in this section. Before we discuss automatic validation, let us discuss XML type modifier.

5.2 XML type modifier

Automatic validation requires XML type modifier.

The XML data type can accept any well-formed XML documents. However, in many cases, users want to store in one XML column documents that have similar structures or conform to the same XML schema. DB2 10 introduces the XML type modifier which qualifies the XML data type with a set of one or more XML schemas. The value of an XML column with an XML type modifier must conform to at least one XML schema specified in the type modifier.

When you define an XML column, you can add an *XML type modifier*. An XML type modifier associates a set of one or more XML schemas with the XML data type. You can use an XML type modifier to cause all XML documents that are stored in an XML column to be validated according to one of the XML schemas that is specified in the type modifier.

The XML type modifier can identify more than one XML schema. You might want to associate more than one XML schema with an XML type modifier for the following reasons:

- ▶ The requirements for an XML schema evolve over time.

An XML column might contain documents that describe only one type of information, but some fields in newer documents might need to be different from fields in the older documents. As new document versions are required, you can add new XML schemas to the XML type modifier.
- ▶ A single XML column contains XML documents of different kinds.

An XML column might contain documents that have several different formats. In this case, each type of document needs its own XML schema.

Alternatively, you might want to associate a single XML schema with multiple type modifiers. An XML schema can define many different documents. You might need to separate the XML

documents into different columns, but specify the same XML schema in a type modifier for each column.

For example, a sales department might have one XML schema that defines purchase orders and billing statements. You can store purchase orders in one XML column, and billing statements in another XML column. Both XML columns have an XML type modifier that points to the same XML schema, but each column restricts documents with different root elements in the XML schema.

You define an XML type modifier in a CREATE TABLE or ALTER TABLE statement as part of an XML column definition.

Not all XML schemas that the XML type modifier identifies need to be registered before you execute the CREATE or ALTER statement. If the XML type modifier specifies a target namespace, only the XML schemas in that target namespace that exist when the CREATE or ALTER statement is executed are associated with the XML type modifier.

If altered-data-type is XML, the old data type of the altered column must also be XML:

- ▶ If the old data type has no XML type modifier and the new data type does, you should ensure that all values in the XML column are valid according to the XML schema that is specified in the type modifier. The XML table space for the column that is being changed is left in CHECK-pending status.
- ▶ If the old data type has the XML type modifier but the new data type has no type modifier, the existing values do not need to be re-validated. The state of the table space is not changed.
- ▶ If the XML schemas that are specified in the old XML type modifier are a subset of the XML schemas that are specified in the new XML type modifier, the existing values do not need to be re-validated. The state of the XML table space is not changed.
- ▶ If the XML schemas that are specified in the old XML type modifier are NOT a subset of the XML schemas that are specified in the new XML type modifier, the XML table space for the column that is being changed is left in the CHECK-pending status.

Changing an XML column to use a different type modifier does not result in the invalidation of dependent plans, packages, or statements in the dynamic statement cache. Also, changing an XML column to use a different type modifier will not generate a new version of the table.

Figure 5-1 shows the XML schemas the following examples refer to for defining an XML type modifier.

XML schema name	Target Namespace	Schema Location	Registration Timestamp
PO1	http://www.example.com/PO1	http://www.example.com/PO1.xsd	2009-01-01 10:00:00.0000
PO2	http://www.example.com/PO2	http://www.example.com/PO2.xsd	2010-01-01 10:00:00.0000
PO3	NO NAMESPACE	http://www.example.com/PO3.xsd	2010-01-30 10:00:00.0000
PO4	http://www.example.com/PO2	http://www.example.com/PO4.xsd	2010-02-23 08:00:00.000

Figure 5-1 XML schemas

Example 5-1 shows how to specify an XML type modifier for an XML column at create time

Example 5-1 Specify an XML type modifier for an XML column at create time

```
CREATE TABLE PURCHASEORDERS (
  ID INT NOT NULL,
  CONTENT XML(XMLSCHEMA ID SYSXSR.PO1))
```

A table for purchase orders contains an XML column named CONTENT. The documents in the XML column need to be validated according to XML schema SYSXSR.PO1, which has already been registered.

To alter an existing XML column to include an XML type modifier or remove an XML type modifier, use ALTER TABLE.

Example 5-2 shows the table definition without XML type modifier specified.

Example 5-2 Table definition without XML type modifier

```
CREATE TABLE PURCHASEORDERS (
  ID INT NOT NULL,
  CONTENT XML)
```

The table contains several XML documents. The documents in the XML column need to be validated according to XML schema SYSXSR.PO1, which has already been registered. Alter the XML column to add an XML type modifier that specifies SYSXSR.PO1 as shown in Example 5-3.

Example 5-3 Specify XML type modifier for XML column at alter time

```
ALTER TABLE PURCHASEORDERS
  ALTER CONTENT
  SET DATA TYPE XML(XMLSCHEMA ID SYSXSR.PO1)
```

Note: The table space that contains the XML documents for the CONTENT column is put in CHECK-pending status.

You can add an XML schema to the XML type modifier.

Example 5-4 shows an example. Suppose PO2 is a new version of the purchase order schema. You can use the ALTER TABLE statement to reset the XML type modifier of the CONTENT column to include both PO1 and PO2.

Example 5-4 Add an XML schema to the XML type modifier

```
ALTER TABLE PURCHASEORDERS
  ALTER CONTENT
  SET DATA TYPE XML(XMLSCHEMA ID SYSXSR.P01, ID SYSXSR.P02)
```

XML schema note: Because the XML schema specified in the old type modifier is a subset of the new type modifier, the existing values of the CONTENT column do not need to be validated again. Thus, the state of the XML table space for the CONTENT column stays unchanged. If the XML schema specified in the old XML type modifier is *not* a subset of the XML schema specified in the new XML type modifier, the XML table space for the column that is being changed is left in the CHECK-pending status.

You can also reset the data type of CONTENT to XML without type modifier.

Example 5-5 shows how to reset XML type modifier for an XML column at alter time.

Example 5-5 Reset XML type modifier for XML column at alter time

```
ALTER TABLE PURCHASEORDERS
  ALTER CONTENT
  SET DATA TYPE XML
```

The CONTENT column in this example: The existing values of the CONTENT column do not need to be validated again.

Validation is automatically performed for INSERT and UPDATE SQL statements and the LOAD utility if the XML column is defined with the XML type modifier.

Catalog table support for XML type modifiers:

- ▶ SYSIBM.SYSXMLTYPMOD contains rows for the XML type modifiers
- ▶ SYSIBM.SYSXMLTYPMSHEMA contains a row for each XML schema specification for one XML type modifier.

Instead of specifying the schema name directly as shown in all the examples, it is possible to specify the URI and LOCATION keywords, so the schema name can be derived.

Example 5-6 shows how to specify the schema location hint. Both PO2 and PO4 have the same target namespace `http://www.example.com/PO2`. If you want to use PO2, you can add LOCATION `'http://www.example.com/PO2.xsd'` after the URI `'http://www.example.com/PO2'` clause.

Example 5-6 Identify an XML schema by target namespace and schema location

```
CREATE TABLE PURCHASEORDERS(
  ID INT NOT NULL,
  CONTENT XML(XMLSCHEMA URI 'http://www.example.com/PO2'
              LOCATION 'http://www.example.com/PO2.xsd' ))
```

Example 5-7 shows XML type modifier uses only the URI keyword to identify the XML schema.

Example 5-7 Identify an XML schema by target namespace

```
CREATE TABLE PURCHASEORDERS(
  ID INT NOT NULL,
  CONTENT XML(XMLSCHEMA URI 'http://www.example.com/PO2'))
```

Using the URI keyword: If you execute the CREATE TABLE statement before PO4 is registered, only PO2 is added to the type modifier in SYSIBM.SYSXMLTYPMSHEMA. When PO4 is registered later, the XML type modifier for the CONTENT column remains unchanged. If you execute the CREATE TABLE statement after PO4 is registered, an SQL error occurs because the XML type modifier uses the URI keyword to identify two XML schemas PO2 and PO4. The URI keyword must identify only one XML schema.

If an XML schema does not contain the targetNamespace attribute in its schema document, it can be referenced in the XML type modifier by "NO NAMESPACE".

In Example 5-8, DB2 chooses PO3 as the XML type modifier.

Example 5-8 No namespace

```
CREATE TABLE PURCHASEORDERS(
  ID INT NOT NULL,
  CONTENT XML(XMLSCHEMA NO NAMESPACE
              LOCATION 'http://www.example.com/PO3.xsd' ))
```

If an XML schema has more than one global element declaration and you want to validate the XML value against one of them, you can specify the ELEMENT clause. Assume the purchase order schema has declared two global elements: purchaseOrder and comment. Hence, a document whose root element is either purchaseOrder or comment could be valid according to PO1 and can be stored in the PURCHASEORDERS table. However, this might not be desirable. If you only want to store purchase order documents in the CONTENT column, you can specify the ELEMENT "purchaseOrder" in the XML type modifier for CONTENT. Example 5-9 shows how you can do this.

Example 5-9 Specifying global element name

```
CREATE TABLE PURCHASEORDERS(
  ID INT NOT NULL,
  CONTENT XML(XMLSCHEMA ID SYSXSR.PO1
              ELEMENT "purchaseOrder"))
```

5.3 Automatic validation

You can automate XML schema validation by adding an XML type modifier to an XML column definition. Before schema validation through an XML type modifier can occur, all schema documents that make up an XML schema must be registered in the built-in XML schema repository (XSR).

Figure 5-2 shows the table definition and the entries in the XSR for the schemas used as XML type modifiers in the table definition.

XML schemas in XML schema repository

XML schema name	Target Namespace	Schema Location	Registration Timestamp
PO1	http://www.example.com/PO1	http://www.example.com/PO1.xsd	2009-01-01 10:00:00.0000
PO2	http://www.example.com/PO2	http://www.example.com/PO2.xsd	2010-01-01 10:00:00.0000
PO3	NO NAMESPACE	http://www.example.com/PO3.xsd	2010-01-30 10:00:00.0000
PO4	http://www.example.com/PO2	http://www.example.com/PO4.xsd	2010-02-23 08:00:00.0000

Table definition:

```
CREATE TABLE PURCHASE_ORDERS(
  ID          INT NOT NULL
  CONTENT XML (XMLSCHEMA ID SYSXSR.PO1, ID SYSXSR.PO2,
              ID SYSXSR.PO3, ID SYSXSR.PO4)
)
```

Figure 5-2 XML Schemas in XML Schema Repository

Figure 5-3 shows how DB2 chooses an XML schema when the XML type modifier specifies multiple schemas.

Schema determination

- DB2 determines the XML schema to use upon INSERT, UPDATE or LOAD.

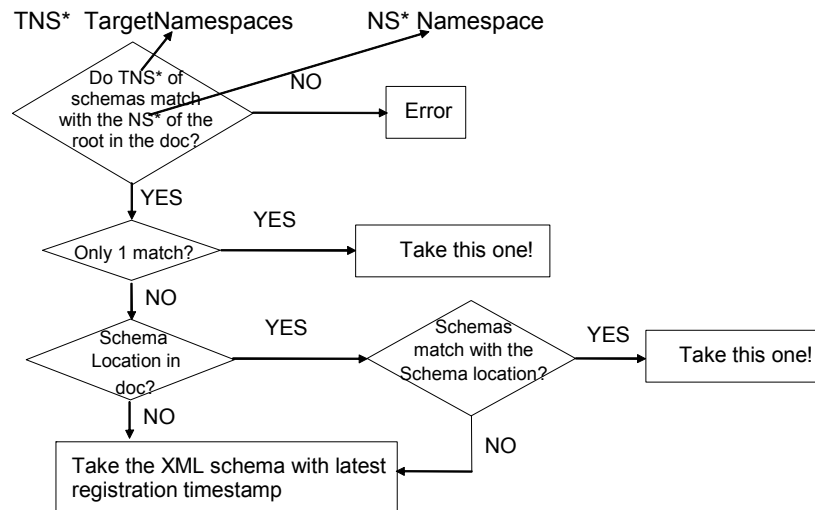


Figure 5-3 Schema determination

You can include more than one XML schema in an XML type modifier. When you insert into or update an XML column, DB2 chooses one XML schema to do validation.

DB2 uses the following process to determine which XML schema to use.

- ▶ If the operation is an update operation, and an XML schema that is specified by the XML type modifier has already been used to validate the original document, DB2 uses the same XML schema to validate the updated document.
- ▶ If there is only one XML schema whose target namespace matches the namespace name of the root element node in the document that is being validated (the XML instance document), DB2 chooses that XML schema to validate the XML document.
- ▶ If there is more than one XML schema with a target namespace that matches the namespace name of the root element, DB2 chooses an XML schema by using the schema location hint. The root element node of an XML instance document can contain an `xsi:schemaLocation` attribute. That attribute consists of one or more pairs of URI references, separated by white space. The first member of each pair is a namespace name, and the second member of the pair is a URI that describes where to find an appropriate schema document for that namespace. The second member of each pair is the schema location hint for the namespace name that is specified in the first member.

For example, this is a schema location attribute:

```
xsi:schemaLocation="http://www.example.com/PO2 http://www.example.com/PO4.xsd"
```

The first member of the pair, `http://www.example.com/PO2`, is the namespace name. The second member of the pair, `http://www.example.com/PO4.xsd`, is the URI that provides the schema location hint.

DB2 uses the schema location hint to choose an XML schema in the following way:

- ▶ If the root element node contains an `xsi:schemaLocation` attribute, DB2 searches the attribute value for a schema location hint with a corresponding namespace name that matches the namespace name in the root element node.
- ▶ If DB2 finds a schema location hint, DB2 uses the hint to identify an XML schema whose schema location URI is identical to the schema location hint. DB2 validates the input document against that schema.
- ▶ If the root element does not contain an `xsi:schemaLocation` attribute, or the `xsi:schemaLocation` attribute does not contain a schema location hint with a corresponding namespace name that matches the namespace name in the root element node, DB2 uses the XML schema with the same target namespace and the latest registration timestamp.

Some examples on how DB2 determines the schema to be used for validation from an XML type modifier are listed here.

In Example 5-10 DB2 chooses XML schema PO1.

Example 5-10 Schema selection for validation from an XML type modifier - Example 1

```
INSERT INTO PURCHASE_ORDERS VALUES(1,
'<po:purchaseOrder xmlns:po="http://www.example.com/PO1">
...
</po:purchaseOrder>');
```

The namespace name in the root element of the instance document is `http://www.example.com/PO1`. This name matches only the target namespace for XML schema PO1.

In Example 5-11 DB2 chooses XML schemas PO2 and PO4 in this order.

Example 5-11 Schema selection for validation from an XML type modifier - Example 2

```
INSERT INTO PURCHASE_ORDERS VALUES (2,
'<po:purchaseOrder xmlns:po="http://www.example.com/PO2"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://www.example.com/PO2
http://www.example.com/PO2.xsd">
...
</po:purchaseOrder>');
```

The namespace name in the root element in the instance document is `http://www.example.com/PO2`, which matches the target namespace of XML schemas PO2 and PO4. The root element of the instance document also contains an `xsi:schemaLocation` attribute whose value provides the schema location hint `http://www.example.com/PO2.xsd`. The schema location hint matches the schema location for XML schema PO2. Therefore DB2 chooses PO2 to validate the instance document. If validation with PO2 fails, DB2 uses PO4.

In Example 5-12 DB2 chooses XML schemas PO4 and PO2 in this order.

Example 5-12 Schema selection for validation from an XML type modifier - Example 3

```
INSERT INTO PURCHASE_ORDERS VALUES (3,
'<po:purchaseOrder xmlns:po="http://www.example.com/PO2"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://www.example.com/PO2
http://www.example.com/PO4.xsd">
...
</po:purchaseOrder>');
```

The namespace name in the root element in the instance document is `http://www.example.com/PO2`, which matches the target namespace of XML schemas PO2 and PO4. The root element of the instance document also contains an `xsi:schemaLocation` attribute whose value provides the schema location hint `http://www.example.com/PO4.xsd`. The schema location hint matches the schema location for XML schema PO4. Therefore DB2 chooses PO4 to validate the instance document. If validation with PO4 fails, DB2 uses PO2.

In Example 5-13 DB2 chooses XML schema PO3.

Example 5-13 Schema selection for validation from an XML type modifier - Example 4

```
INSERT INTO PURCHASE_ORDERS VALUES (4,
'<purchaseOrder xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="http://www.example.com/PO3.xsd">
...
</purchaseOrder>');
```

The root element of the instance document has no namespace name. XML schema PO3 has no target namespace. Therefore, DB2 uses PO3 for validation.

Note that, after you update an XML document in a column that has an XML type modifier, DB2 validates again all or part of the document.

- ▶ If the XML type modifier includes several XML schemas, DB2 uses the same XML schema for validation that is used for the original validation.
- ▶ If you update an entire document, DB2 validates the entire document. However, if you use the `XMLMODIFY`¹ function to update only a portion of the document, DB2 might need to validate only the updated portion.

5.4 User-controlled validation

One other way to do XML schema validation is by executing the SYSIBM.DSN_XMLVALIDATE built-in function. Before you can invoke DSN_XMLVALIDATE, all schema documents that make up an XML schema must be registered in the XML schema repository and successfully compiled.

You can use DSN_XMLVALIDATE with a type modifier. This can be used to override a schema picked by DB2. DB2 checks if the schema used in DSN_XMLVALIDATE is one of the schemas in the type modifier, and skips double validation. There are a number of forms of DSN_XMLVALIDATE:

```
DSN_XMLVALIDATE(string-expression)
DSN_XMLVALIDATE(xml-expression)
DSN_XMLVALIDATE(string-expression, varchar-expression)
DSN_XMLVALIDATE(xml-expression, varchar-expression)
DSN_XMLVALIDATE(string-expression1, string-expression2, string-expression3)
DSN_XMLVALIDATE(xml-expression1, string-expression2, string-expression3)
```

For all forms, the first parameter contains the document that you want to validate.

For forms with one parameter, the target namespace and optional schema location of the XML schema must be in the root element of the instance document that you want to validate.

For forms with two parameters, the second parameter is the name of the schema object to use for validation of the document. That object must be registered in the XML schema repository.

For forms with three parameters, the second and third parameter contain the names of a namespace URI and a schema location hint that identify the XML schema object to use for validation of the document. That object must be registered in the XML schema repository

XML schemas used for validation are the same as shown in Figure 5-2 on page 79. However, the following CREATE TABLE statement is used to create the table:

```
CREATE TABLE PURCHASE_ORDERS(
  ID      INT      NOT NULL
  CONTENT XML )
```

Examples of the criteria how DB2 chooses XML schema for DSN_XMLVALIDATE are listed here.

In Example 5-14 DB2 chooses XML schema PO1.

Example 5-14 Schema selection for validation for DSN_XMLVALIDATE - Example 1

```
INSERT INTO PURCHASE_ORDERS VALUES(1,
DSN_XMLVALIDATE('<po:purchaseOrder xmlns:po="http://www.example.com/PO1">
...
</po:purchaseOrder>'));
```

The DSN_XMLVALIDATE invocation does not specify an XML schema or target namespace and schema location hint, so DB2 uses the information in the instance document.

¹ The XMLMODIFY function returns an XML value that might have been modified by the evaluation of an XPath updating expression and XPath variables that are specified as input arguments.

The namespace name in the root element of the instance document is `http://www.example.com/PO1`. This name matches only the target namespace for XML schema PO1.

In Example 5-15 DB2 chooses XML schema PO2.

Example 5-15 Schema selection for validation for DSN_XMLVALIDATE - Example 2

```
INSERT INTO PURCHASE_ORDERS VALUES (2,
DSN_XMLVALIDATE('<po:purchaseOrder xmlns:po="http://www.example.com/PO2"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://www.example.com/PO2
http://www.example.com/PO2.xsd">
...
</po:purchaseOrder>', 'SYSXSR.PO2'));
```

The `DSN_XMLVALIDATE` invocation specifies XML schema `SYSXSR.PO2`.

In Example 5-16 DB2 chooses XML schema PO4.

Example 5-16 Schema selection for validation for DSN_XMLVALIDATE - Example 3

```
INSERT INTO PURCHASE_ORDERS VALUES (3,
DSN_XMLVALIDATE('<po:purchaseOrder xmlns:po="http://www.example.com/PO2"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://www.example.com/PO2
http://www.example.com/PO4.xsd">
...
</po:purchaseOrder>', 'http://www.example.com/PO2'));
```

The `DSN_XMLVALIDATE` invocation specifies namespace `http://www.example.com/PO2`. Two XML schemas, PO2 and PO4, have that target namespace. DB2 uses PO4, because it has the later timestamp.

In Example 5-17 DB2 chooses PO3.

Example 5-17 Schema selection for validation for DSN_XMLVALIDATE - Example 4

```
INSERT INTO PURCHASE_ORDERS VALUES (4,
DSN_XMLVALIDATE('<purchaseOrder
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="http://www.example.com/PO3.xsd">
...
</purchaseOrder>');
```

The `DSN_XMLVALIDATE` invocation does not specify an XML schema or target namespace and schema location hint, so DB2 uses the information in the instance document. The root element node in the instance document contains an `xsi:noNamespaceSchemaLocation` attribute with value `http://www.example.com/PO3.xsd`, so DB2 uses XML schema PO3, which has no target namespace, and the schema location `http://www.example.com/PO3.xsd`.

There have been two versions of `DSN_XMLVALIDATE`:

- ▶ A user-defined function
- ▶ A built-in function

The user-defined function is deprecated. Now DB2 uses the built-in function instead even in DB2 9.

To move from the DSN_XMLVALIDATE user-defined function to the DSN_XMLVALIDATE built-in function:

- ▶ For applications that invoke DSN_XMLVALIDATE using the qualified name SYSFUN.DSN_XMLVALIDATE:
 - a. Change the name to SYSIBM.DSN_XMLVALIDATE.

This change is optional. DB2 drops SYSFUN.DSN_XMLVALIDATE and invalidates packages during migration. Automatic rebinds pick up SYSIBM.DSN_XMLVALIDATE.
 - b. Prepare the applications again.
- ▶ For applications that invoke DSN_XMLVALIDATE without using the qualified name, you do not need to modify the applications. DB2 uses the SYSIBM.DSN_XMLVALIDATE built-in function automatically.
- ▶ Optional: Remove the XMLPARSE function that surrounds DSN_XMLVALIDATE.

The SYSFUN.DSN_XMLVALIDATE user-defined function must be invoked from within the XMLPARSE function. The SYSIBM.DSN_XMLVALIDATE built-in function does not need to be invoked from within the XMLPARSE function.

Deprecated function: The SYSFUN.DSN_XMLVALIDATE user-defined function is deprecated. Use the SYSIBM.DSN_XMLVALIDATE built-in function instead. Even if you explicitly call SYSFUN.DSN_XMLVALIDATE, DB2 runs SYSIBM.DSN_XMLVALIDATE.

5.5 Determining whether an XML document has been validated

You can use the SQL XMLXSROBJECTID scalar function to determine whether an XML document that is stored in a table has undergone XML validation, and which XML schema was used to validate that document.

XMLXSROBJECTID returns the XSR object identifier of the XML schema that was used to validate the input XML document. The XSR object identifier corresponds to the XSROBJECTID column in the SYSIBM.XSROBJECTS “catalog” table. After you call XMLXSROBJECTID, you can use the returned value to query SYSIBM.XSROBJECTS for the XML schema information. If the XML document has not been validated, XMLXSROBJECTID returns 0.

The SQL statement in Example 5-18 calls XMLXSROBJECTID to determine which XML documents in the INFO column of the CUSTOMER table have not been validated. The statement then calls DSN_XMLVALIDATE to validate those documents against XML schema SYSXSR.P01.

Example 5-18 Search for documents not validated

```
UPDATE CUSTOMER
SET INFO = DSN_XMLVALIDATE(INFO, 'SYSXSR.P01')
WHERE XMLXSROBJECTID(INFO)=0
```

The SQL statement in Example 5-19 retrieves the XML schema names and target namespaces for the XML schemas that were used to validate XML documents in the INFO column of the CUSTOMER table.

Example 5-19 Retrieve target namespaces and XML schema names used for validation

```
SELECT DISTINCT S.XSROBJECTNAME, S.TARGETNAMESPACE  
FROM CUSTOMER C, SYSIBM.XSROBJECTS S  
WHERE XMLXSROBJECTID(INFO) = S.XSROBJECTID
```



DB2 SQL/XML programming

In this chapter, we provide a wide range of DB2 programming examples for pureXML. Over recent years the amount of programming that is deployed inside DB2 has been growing with the adoption of facilities like stored procedures, user defined functions, triggers and WebSphere MQ integration.

The adoption of pureXML is likely to increase this trend. The ease of handling XML documents in native SQL procedures, compared to external language environments like COBOL, means that development productivity will be enhanced by encapsulating XML logic within DB2 procedures and functions, so that external programs only have to deal with traditional SQL-based functions.

This chapter addresses the following DB2 programming topics:

- ▶ Native SQL stored procedures and XML
- ▶ Receiving XML messages from MQ
- ▶ Audit queries (against logged XML messages)
- ▶ SQL/XML query techniques
- ▶ User defined functions with XML
- ▶ Triggers with XML
- ▶ XML joins
- ▶ XML with change data capture tools

6.1 Native SQL stored procedures and XML

The book application scenario, XML message logging and auditing, described in Chapter 3, “Application scenario” on page 45, represents a common situation, where XML messages are flowing over a WebSphere MQ enterprise service bus, and you would like to capture some of this messages and store them in DB2 pureXML for auditing purposes.

Native stored procedures are an excellent vehicle for capturing XML messages from messaging infrastructure (like WebSphere MQ and DataPower®) because they are:

- ▶ Productive: they allow you to encapsulate multi-step processes into a single DB2 callable routine, which supports the XML data type explicitly.
- ▶ Well suited to handling XML: they can use XML documents as input and output parameters, and manipulate XML without necessarily having to persist it to a DB2 table.
- ▶ Efficient: they can parse the incoming XML document once, and re-use it as an XML data type without re-parsing
- ▶ Well connected: they integrate well with MQ

Native stored procedures become more powerful as a result of their support for XML. Prior to DB2 10, a DB2 stored procedure could only be passed individual data elements as parameters. Now, it is possible to pass them arrays of data within an XML input parameter, making it much more practical to implement large scale processing work within the DB2 subsystem, where it is most efficient.

This section shows examples of XML handling with native SQL stored procedures, and then adds integration with WebSphere MQ.

We use the table definitions shown in Example 6-1 for the initial examples. The sample stored procedures show techniques to validate the data manually, as well as automatically, which is why we have two versions of the audit logging table.

Example 6-1 Tables used for following examples.

```
-- create base table for storing ISO20022 XML documents, with validation

CREATE TABLE BK_TO_CSTMRT_STMT (
  MSG_ID VARCHAR(35) ,
  MSG_CRE_DT_TM TIMESTAMP,
  BK_TO_CSTMRT_STMT XML(XMLSCHEMA ID SYSXSR.SG247915_01) NOT NULL)
IN XMLR3DB.TSAUDIT1 ;

-- create alternate base table for ISO20022 XML documents, without validation

CREATE TABLE BK_TO_CSTMRT_STMT_MANUALVALIDATE (
  MSG_ID VARCHAR(35) ,
  MSG_CRE_DT_TM TIMESTAMP,
  BK_TO_CSTMRT_STMT XML NOT NULL)
IN XMLR3DB.TSAUDIT2 ;

-- create iso20022 currency lookup table

CREATE TABLE CURRENCY (
  ENTITY VARCHAR(50),
  CURRENCY VARCHAR(50),
  ALPHABETIC_CODE CHAR(3),
```



```

    NUMERIC_CODE SMALLINT )
  IN XMLR3DB.TSAUDIT3 ;

-- create table for error handling logic to store invalid documents

CREATE TABLE XMLR3.INVALID_DOCS (
  INSERT_TIME TIMESTAMP,
  INSERT_DOC XML,
  ERROR_MESSAGE VARCHAR(1000) )

```

We are also using the ISO20022 bank-to-customer-statement schema, which we have taken directly from the ISO20022 web site without any alteration. We have downloaded the schema to D:\XMLRES\WEEK3\REDXMPL\camt.053.001.02.xsd and registered the schema with name SYSXSR.SG247915_01 in the DB2 XSR, using the commands in Example 6-2. Note that the commands to register schemas in the XSR are available in two environments: z/OS UNIX System Services, and DB2 for Linux, UNIX and Windows Command Line Processor (CLP). They are explained in more detail in 2.1.6, “XML schema repository and schema validation” on page 36.

Example 6-2 Registering the ISO20022 Bnk_To_Cst_Stmt XML schema

```

register xmlschema 'redbook_bankstmt.xsd' from
  file://D:\SG247915\camt.053.001.02.xsd as SYSXSR.SG247915_01 ;

DB20000I The REGISTER XMLSCHEMA command completed successfully.

complete xmlschema SYSXSR.SG247915_01 ;

DB20000I The COMPLETE XMLSCHEMA command completed successfully.

```

Alternatively, in Example 6-2, the *register xmlschema* command can have the *complete* clause at the end to complete the registration in the same command.

6.1.1 Native SQL stored procedure example

Let us first show an example of how XML can be used with a native SQL stored procedure. The sample procedure in Example 6-3 has been written to:

1. Receive an XML document as an input parameter
2. Validate the incoming XML document against the schema in DB2 XSR
3. Use various XML functions to extract elements from the incoming XML document
4. Store the document (and extracted fields) in the DB2 audit logging table
5. Use XML publishing functions to generate a new XML document
6. Return the generated XML document as an output parameter

Example 6-3 Simple stored procedure for registering the ISO20022 Bnk_To_Cst_Stmt XML schema

```

CREATE PROCEDURE STOREXML1 (
  IN V_BANKSTMT XML,
  OUT V_MSG_ID VARCHAR(35),
  OUT V_CREDITM TIMESTAMP,

```

```

OUT V_MINISTMT XML) ❶
LANGUAGE SQL
MODIFIES SQL DATA
DISABLE DEBUG MODE

BEGIN

DECLARE VALIDXML XML ;
DECLARE SQLCODE INTEGER ;

SET VALIDXML = DSN_XMLVALIDATE(V_BANKSTMT, 'SYSXSR.SG247915_01') ; ❷

SET V_MSG_ID = (
  xmlcast(xmlquery(
    'declare default element
     namespace"urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $d/Document/BkToCstmrStmt/GrpHdr/MsgId'
    passing VALIDXML as "d")
  as varchar(35))); ❸

SET V_CREDITM = (
  xmlcast(xmlquery(
    'declare default element namespace
     "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $d/Document/BkToCstmrStmt/GrpHdr/CredtTm'
    passing VALIDXML as "d")
  as timestamp)); ❹

SET V_MINISTMT = (
  xmlquery('declare default element namespace
     "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $d/Document/BkToCstmrStmt/Stmt'
    passing VALIDXML as "d")); ❺

INSERT INTO BK_TO_CSTMRS_STMT_MANUALVALIDATE (
  MSG_ID , MSG_CRE_DT_TM, BK_TO_CSTMRS_STMT)
  values ( V_MSG_ID, V_CREDITM, VALIDXML ) ; ❻

END !

```

The following notes explain each of the annotated points in the source code for the stored procedure.

1. The parameter list includes one input parameter (data type XML) and three output parameters (data types VARCHAR(35), Timestamp, and XML). The ability to support XML as parameters to a callable routine is new in DB2 10.
2. This statement uses the system-provided DSN_XMLVALIDATE function to validate the input XML document against a schema that is defined in the DB2 XSR, WITHOUT storing the XML document in a DB2 table. (Most use cases show schema validation as part of an SQL insert or update).
3. This statement shows the use of the XMLQUERY function to strip a string data element from the XML document, and then cast it to a relational data type using XMLCAST.

4. This statement shows another example of the XMLQUERY function to strip a different field with data type timestamp from the XML document. (DB2 10 supports both timestamp and timestamp with timezone data types in both relational and XML structures).
5. This statement uses an XMLQUERY function to generate a new XML document from the received XML document, and to include the generated XML document as an output parameter from the stored procedure.
6. This statement stores the received XML document, alongside two stripped data elements in a DB2 table.

One of the strengths of the native SQL stored procedure is the ability to operate on an XML document in memory, without storing it in a DB2 table (as performed in steps 2, 3, 4 and 5). This ability makes the stored procedure efficient because it saves making unnecessary WRITE and READ operations to a DB2 table.

This stored procedure contains XML data types which make it difficult to test from many of the tools which you may be using, because they do not have an easy way to pass XML types to & from the procedure (such as: SPUFI, DB2 Command Line Processor etc...). A Java program was used as a test driver to call the stored procedure from a windows db2 client: it calls the stored procedure with an XML input parameter, and receives the output parameters, and writes them to an output file. This program is called Teststorexml1.java, and is included in the additional materials of this book.

This stored procedure example could be improved in several ways, which we develop in this chapter. For example:

- ▶ No error handling logic is included within the stored procedure.
- ▶ The three XMLQUERY operations could be replaced by a single XMLTABLE operation

6.1.2 XML error handling in native SQL procedures

The principles of error handling for XML processing is no different from normal native SQL stored procedures. There are many SQLCODES and SQLSTATES that provide diagnostic information in the event of an error relating to the various XML related functions and procedures in DB2. These error conditions can be handled in the same way as relational errors. You just need to be aware of new situations where errors might arise, and code for them.

The basic approach to error handling in native stored procedures is to:

1. Anticipate which error conditions are likely to be encountered during normal use of the system you are building. For example, if you are receiving XML documents from external sources, you must consider that some may be badly formed or fail XML schema validation.
2. Define error conditions that are likely to arise
3. Define error handling routines that take the appropriate action for an anticipated circumstance (and a catch-all routine for any other error). For example, if a received XML document fails validation, you may want to store it as a CLOB in a separate table for inspection.

The previous SQL stored procedure in Example 6-3 has been modified to include an XMLTABLE function (to replace the two separate XMLQUERY functions) and to include some basic error handling. The modified procedure is listed in Example 6-4.

Example 6-4 Stored procedure with error handling logic

```
CREATE PROCEDURE STOREXML2 (
  IN V_BANKSTMT XML,
```

```

OUT V_MSG_ID VARCHAR(35),
OUT V_CREDITM TIMESTAMP,
OUT V_MINISTMT XML,
OUT ERROR_MESSAGE VARCHAR(1000) )
LANGUAGE SQL
MODIFIES SQL DATA
DISABLE DEBUG MODE

BEGIN

DECLARE v_sql VARCHAR(2048) ;
DECLARE VALIDXML XML ;
DECLARE SQLCODE INTEGER ;
DECLARE SQLSTATE CHAR(5) ;
DECLARE SQLERRMC VARCHAR(70) ;
DECLARE ERROR_MESSAGE VARCHAR(250) ;
DECLARE INVALID_DOCUMENT CONDITION FOR SQLSTATE '2200M'; ❶

DECLARE c1 CURSOR with return to caller FOR stmt;

DECLARE EXIT HANDLER FOR INVALID_DOCUMENT
BEGIN
set ERROR_MESSAGE = 'DB2 DIAGNOSTICS - SQLCODE= '
concat CHAR(SQLCODE)
concat ' SQLSTATE= '
concat SQLSTATE
concat ' SQLERRMC= '
concat SQLERRMC ;
insert into INVALID_DOCS ( INSERT_TIME, INSERT_DOC, ERROR_MESSAGE )
values ( current timestamp, V_BANKSTMT, ERROR_MESSAGE ) ;
END ; ❷

SET VALIDXML = (
DSN_XMLVALIDATE(V_BANKSTMT, 'SYSXSR.SG247915_01')) ; ❸

SELECT X.MSG_ID, X.CRE_DT_TM, X_MINISTMT
INTO V_MSG_ID, V_CREDITM, V_MINISTMT
FROM XMLTable(XMLNAMESPACES(
DEFAULT 'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02'),
'$d/Document/BkToCstmrStmt' PASSING VALIDXML as "d"
COLUMNS
"MSG_ID" VARCHAR(35) PATH './GrpHdr/MsgId',
"CRE_DT_TM" TIMESTAMP PATH './GrpHdr/CreDtTm',
"X_MINISTMT" XML PATH './Stmt' ) AS X ; ❹

INSERT INTO BK_TO_CSTMTR_STMT_MANUALVALIDATE
(MSG_ID , MSG_CRE_DT_TM, BK_TO_CSTMTR_STMT)
values ( V_MSG_ID, V_CREDITM, VALIDXML ) ;

END !

```

The following notes explain each of the annotated points in the source code for the stored procedure.

1. This statement declares an anticipated error condition for the validation of the received XML document.
2. This statement defines the processing logic when the declared error conditions are encountered.
3. This call to the DSN_XMLVALIDATE function is the call that is likely to generate the error condition that the error handling routine is setup for.
4. This statement is an XMLTABLE operation, which is more efficient than multiple XMLQUERY operations when you want to strip multiple elements from the XML document.

Another Java program (Teststorexml2) is included in the additional materials of this book to act as a test driver for the stored procedure.

6.1.3 Stored procedures development tools

The focus of this book is the XML capabilities of DB2 for z/OS. However, it is important to be also aware of the development tools that are available for building native stored procedures with XML.

Simple native SQL procedures, like the ones above, can easily be coded and tested with a text editor. More complex stored procedures will benefit from the stored procedure development and debug capabilities of IBM Data Studio.

Data Studio can be downloaded, free of charge, from

<http://www.ibm.com/developerworks/downloads/im/data/>

Data Studio provides a wide range of support for developers and DBA. Specifically, with regard to SQL stored procedures, Data Studio

- ▶ Supports the build, test, optimization and deployment of SQL stored procedures with an interactive routine debugger.
- ▶ Provides drag and drop query builders with editors for SQL and SQL/XML.
- ▶ Supports XML data type too.

For detailed documentation on using Data Studio for stored procedure development refer to Part 6. “Cool tools for an easier life” of *DB2 9 for z/OS Stored Procedures: Through the CALL and Beyond*, SG24-7604.

6.2 Receiving XML messages from MQ

The application scenario in Chapter 3, “Application scenario” on page 45 describes the increasingly common environment where organizations make extensive use of event driven, messaging and workload systems. It is very common that such systems will send XML messages over their Enterprise Service Bus.

In enterprises that use z/OS mainframes, it is very common that WebSphere MQ is the transport for XML messages. DB2 provides a range of DB2 scalar functions and table functions for working with WebSphere MQ. DB2 also provides the MQ listener, which receives messages from a message queue and calls a DB2 stored procedure with the contents of the message.

This section provides worked examples of how the ISO20022 Bank-To-Customer-Statement message can be received from WebSphere MQ, and ingested by DB2.

6.2.1 WebSphere MQ functions

DB2 provides a range of built in functions for integrating with WebSphere MQ. These functions have been updated in DB2 10, so that they are based on the newer MQI interface to WebSphere MQ, rather than the older AMI interface which was used in DB2 8.

- ▶ Installation job DSNTIJRT should have been run to install the DB2 MQ functions.
- ▶ Installation job DSNTJMQ verifies the MQ environment setup.

The MQ functions in DB2 10 reside in schema DB2MQ and run in WLM environment DSNWLM_MQSERIES.

In order to use the DB2 MQ functions within SQL statements, the MQ functions must be defined as services to DB2. These service definitions are used to encapsulate the MQ programming properties that MQ requires, so that the DB2 programmer can use a range of relatively simple functions to access WebSphere MQ. All the programmer has to do is use a system provided MQ function (such as MQSEND), in conjunction with an MQ service name.

The MQ services and policies must be defined in the following tables:

- ▶ SYSIBM.MQSERVICE_TABLE
 - It contains a list of "MQ services" available within a DB2 subsystem.
 - An MQ service is a definition of a Queue, its Queue Manager, and codepage properties of that queue.
- ▶ SYSIBM.MQPOLICY_TABLE
 - It contains a list of "MQ service policies" that can be used.
 - An MQ service policy defines MQ properties, such as message priorities, retries, exception handling that are to be used by DB2 when accessing a message queue.

We have defined an MQ service using the SQL insert in Example 6-5. We have not explicitly defined an MQ service policy, which means that we can call the DB2 MQ functions without specifying an MQ service policy, and we implicitly accept the default MQ service policy that was defined by the installation jobs.

Example 6-5 Populating the MQSERVICE-TABLE

```
INSERT INTO SYSIBM.MQSERVICE_TABLE (
  SERVICENAME, QUEUEMANAGER, INPUTQUEUE, CODEDCHARSETID, ENCODING, DESCRIPTION)
VALUES ('XMLS1', 'MQBA', 'MQL.INPUT01', DEFAULT, DEFAULT, DEFAULT)
```

Actually, it is more efficient to receive an XML document from an external source as a VARCHAR or a CLOB if you are going to validate the document against an XML schema.

The MQ scalar functions provided by DB2 are summarized in Table 6-1.

Table 6-1 MQ scalar functions provided by DB2 10.

DB2 MQ function	Invocation parameters	Description
MQREAD	receive-service, service-policy	It reads a message as a VARCHAR (32704 bytes max) from a specified queue, using a specified policy, and leaves the message at the head of the queue. (empty queue returns null).

DB2 MQ function	Invocation parameters	Description
MQREADCLOB	receive-service, service-policy	It reads a message as a CLOB (2 MB max) from a specified queue, using a specified policy, and leaves the message at the head of the queue. (empty queue returns null).
MQRECEIVE	receive-service, service-policy, correlation-id	It reads a message (with matching correlation-id, if specified) as a VARCHAR (32704 bytes max) from a specified queue, using a specified policy, and removes the message from the head of the queue. (empty queue returns null).
MQRECEIVECLOB	receive-service, service-policy, correlation-id	It reads a message (with matching correlation-id, if specified) as a CLOB (2 MB max) from a specified queue, using a specified policy, and removes the message from the head of the queue. (empty queue returns null).
MQSEND	send-service, service-policy, msg-data, correlation-id	It writes a message (with correlation-id, if specified) as a VARCHAR (32707 bytes max) or CLOB (2 MB max) to a specified queue, using a specified policy.

DB2 provides additional MQ table functions and MQ publishing functions as listed in the *DB2 10 for z/OS Installation and Migration Guide*, GC219-2974.

Examples of the MQREAD and MQSEND functions that can be executed from SPUFI (or any other SQL editor) are shown in Example 6-6. Note that some of the invocation parameters are optional.

Example 6-6 Sample MQREAD and MQSEND function calls

```
select db2mq.mqread('XMLS1')
  from sysibm.sysdummy1

-- this SQL statement returns the contents of the first message on the queue
-- referenced by the XMLS1 MQ service as a string.

select db2mq.mqsend('XMLS1', '<testxml><tag1>newvalue</tag1></testxml>')
  from sysibm.sysdummy1

-- this SQL statement puts the simple XML string onto the queue referenced by the
-- XMLS1 MQ service as a string.
```

There are no DB2 MQ functions that directly use XML data type parameters. When using functions like MQSEND and MQREAD to send and receive XML documents to and from MQ, the XML document must be passed as a string data type, and then converted to XML.

XML schema validation (where automatic or using DSN_XMLVALIDATE) only accepts input in string format. If you use XML type, DB2 must parse the XML document into internal format during the parameter passing phase, then implicitly serialize it back to string before validation, which would be slower.

6.2.2 DB2 stored procedure reading from MQ

Having understood the MQ scalar functions that are available, it is now possible to write a native stored procedure that reads an XML message from a message queue (as a string), and stores it in a DB2 table (as an XML data type), and performs any other useful processing

against the parsed XML document whilst it is in memory. The stored procedure in Example 6-7 performs the following steps:

1. Reads XML document as a string from WebSphere MQ.
2. Parses the string into an XML data type (an automatically validates that it is well formed).
3. Performs schema validation against a stored XML schema in the XML schema repository
4. Extracts specific fields from the XML document, which we want to store in relational columns in the message logging table.
5. Writes the XML document, and stripped relational fields to the message logging table.

Example 6-7 Stored procedure to read XML message from MQ

```

CREATE PROCEDURE STOREXML3 (
  OUT V_MSG_ID VARCHAR(35),
  OUT V_CREDTTM TIMESTAMP,
  OUT MYOUTPUT VARCHAR(1))
  LANGUAGE SQL
  MODIFIES SQL DATA
  DISABLE DEBUG MODE

BEGIN
DECLARE OUTPUTMQ CLOB ;
DECLARE VALIDXML XML ;
DECLARE MQSVC CHAR(5) ;
DECLARE SQLCODE INTEGER ;
SET MQSVC = 'XMLS1' ;

SET OUTPUTMQ = ( select db2mq.mqread(MQSVC) from sysibm.sysdummy1 ) ; 1

SET VALIDXML = (
  SELECT DSN_XMLVALIDATE(OUTPUTMQ, 'SYSXSR.SG247915_01')
  FROM SYSIBM.SYSDUMMY1 ) ; 2

SELECT X.MSG_ID, X.CRE_DT_TM
  INTO V_MSG_ID, V_CREDTTM
  FROM XMLTable(XMLNAMESPACES(
    DEFAULT 'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02'),
    '$d/Document/BkToCstmrStmt' PASSING VALIDXML as "d"
  COLUMNS
    "MSG_ID" VARCHAR(35) PATH './GrpHdr/MsgId',
    "CRE_DT_TM" TIMESTAMP PATH './GrpHdr/CreDtTm' ) AS X ; 3

INSERT INTO BK_TO_CSTMTR_STMT_MANUALVALIDATE (
  MSG_ID , MSG_CRE_DT_TM, BK_TO_CSTMTR_STMT)
  values ( V_MSG_ID, V_CREDTTM, VALIDXML ) ; 4

END !

```

The numbered highlighted steps in the stored procedure are:

1. This statement reads the XML message from MQ, and assigns it to variable OUTPUTMQ, which is defined as a CLOB.
2. This statement validates the CLOB that was read from WebSphere MQ against a the ISO20022 schema that was registered in DB2 XSR.

3. This statements strips selected data elements from the XML document
4. This statement inserts the XML document and stripped elements into a relational table.

The stored procedure illustrated above is written to read the first message only from the message queue, and process it. In practice you would modify the stored procedure to call the DB2MQ.MQRECEIVE function, so as to receive the message from the queue, and remove it within the unit of work controlled by the DB2 stored procedure.

6.2.3 DB2 MQ Listener automation

The stored procedure in the previous section could be extended to drain the queue of all input messages with a simple programming loop, and the procedure could be scheduled by some application environment to execute periodically. However, a simpler approach would be to use the MQ listener, which will take care of all the application programming and scheduling effort for you. The MQ listener provides the capability to listen to a message queue for messages when they arrive, and automatically call a stored procedure when they do. It saves you writing and scheduling an application to poll MQ for messages and process them when they do arrive.

The MQ Listener is a standard component of DB2, which is installed with job SDSNSAMP(DSNTIJML). The listener runs under USS (unix system services) in z/OS. It can be invoked using USS commands, as shown in the examples that follow. There is a JCL sample to invoke these uss commands, found in SDSNSAMP(DSNTEJML)

Once the MQ Listener has been installed, it must be configured and started. The USS command to configure the MQ Listener to listen on an input queue, is shown in Example 6-8. In this example, we are defining a 2-phase commit MQ Listener process (db2mqln2 command implies this) in DB2 subsystem DB0B to listen to the queue MQL.INPUT01 in Queue Manager MQBA and to call procedure XMLR3.STOREXML4 when a message is received. The configuration is stored under the label XML1.

Example 6-8 Command to configure MQ listener

```
db2mqln2 add
  -ssID DB0B
  -config XML1
  -queueManager MQBA
  -inputQueue MQL.INPUT01
  -procName STOREXML4
  -procSchema XMLR3
  -numInstances 1
```

The configuration can be verified with the show command as shown in Example 6-9.

Example 6-9 Command to show MQ listener configuration

```
db2mqln2 show -ssID DB0B -config all
```

returns:

```
configurationName: XML1
queueManager:      MQBA
inputQueue:        MQL.INPUT01
procSchema:        XMLR3
procName:          STOREXML4
```

```

numInstances:      1
mqCoordinated:    T

```

The underlying DB2 table that stores the MQ Listener configuration data is SYSMQL.LISTENERS, whose contents are shown in Example 6-10.

Example 6-10 Contents of SYSMQL.LISTENERS

SYSMQL.LISTENERS

```

CONFIGURATIONNAME = XML1
QUEUEMANAGER      = MQBA
INPUTQUEUE        = MQL.INPUT01
PROCNODE          =
PROCSHEMA         = XMLR3
PROCNAME          = STOREXML4
PROCTYPE          = 1
NUMINSTANCES     = 1
WAITMILLIS        = 50
MINQUEUEDEPTH    = 1

```

In order for a stored procedure to work with the MQ Listener, it must be coded with INPUT and OUTPUT characters that may only be VARCHAR, VARBINARY, BLOB or CLOB. The stored procedure to read an XML message from a message queue listed in Example 6-7 on page 96 can be modified to work with the MQ Listener by making a few small edits as shown in Example 6-11.

Example 6-11 Stored procedure modified for MQ listener integration

```

CREATE PROCEDURE STOREXML3 (
  OUT V_MSG_ID VARCHAR(35),
  OUT V_CREDITM TIMESTAMP,
  OUT MYOUTPUT VARCHAR(1))
  LANGUAGE SQL
  MODIFIES SQL DATA
  DISABLE DEBUG MODE

BEGIN
  DECLARE OUTPUTMQ CLOB ;
  DECLARE VALIDXML XML ;
  DECLARE MQSVC CHAR(5) ;
  DECLARE SQLCODE INTEGER ;

  SET MQSVC = 'XMLS1' ;

  SET OUTPUTMQ = ( select db2mq.mqread(MQSVC) from sysibm.sysdummy1 ) ;

  SET VALIDXML = ( SELECT DSN_XMLVALIDATE(OUTPUTMQ, 'SYSXSR.SG247915_01')
    FROM SYSIBM.SYSDUMMY1 ) ;

  SELECT X.MSG_ID, X.CRE_DT_TM INTO V_MSG_ID, V_CREDITM
    FROM XMLTable(XMLNAMESPACES(DEFAULT
      'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02'),
      '$d/Document/BkToCstmrStmt' PASSING VALIDXML as "d"
    COLUMNS

```

```

        "MSG_ID" VARCHAR(35) PATH './GrpHdr/MsgId',
        "CRE_DT_TM" TIMESTAMP PATH './GrpHdr/CreDtTm' ) AS X ;
INSERT INTO BK_TO_CSTMTR_STMT_MANUALVALIDATE (
    MSG_ID , MSG_CRE_DT_TM, BK_TO_CSTMTR_STMT)
VALUES ( V_MSG_ID, V_CREDTTM, VALIDXML ) ;

END !

```

The listener can be operated with the commands as shown in Example 6-12.

Example 6-12 Commands to operate MQ listener

```

db2mq1n2 run
  -ssID DB0B
  -config XML1
  -adminQueue MQL.ADMIN
  - adminQMgr MQBA

db2mq1n2 admin
  -adminQueue MQL.INPUT01
  -adminQMgr MQBA
  -adminCommand shutdown

```

These examples have show how native stored procedures can be integrated with WebSphere MQ to receive XML messages, process them and log them to DB2.

6.3 Audit queries (against logged XML messages)

Now that we have received a stream of XML messages from WebSphere MQ, and stored them in DB2, we have the ability to query them.

One of the first things to note is that you don't necessarily need any special "XML-enabled" query tools to develop and run audit queries, because the SQL/XML language provides a range of functions to search within the XML documents for data of interest, without materializing XML structures that the query tool has to handle.

If you do have an XML-enabled query tool, then SQL/XML can of course return the results of queries as XML documents. The examples that follow will include a mixture of both.

The first few examples in this section will include screen shots from Optim Development Studio. This is one of the tool choices outlined in 2.3.2, "GUI based tools" on page 43, and it should provide a flavour for the interface style of the ODS tool.

6.3.1 Simple SQL/XML search examples

The SQL/XML queries in this section will focus on the very simple case of our BK_TO_CSTMTR_STMT table with 5 rows populated in it. The 5 rows represent 5 account statements from January 2010 through to May 2010.

First, let us look at the entire table. Figure 6-1 shows the Optim Development Studio being used to execute the "SELECT * FROM BK_TO_CSTMTR_STMT" statement.

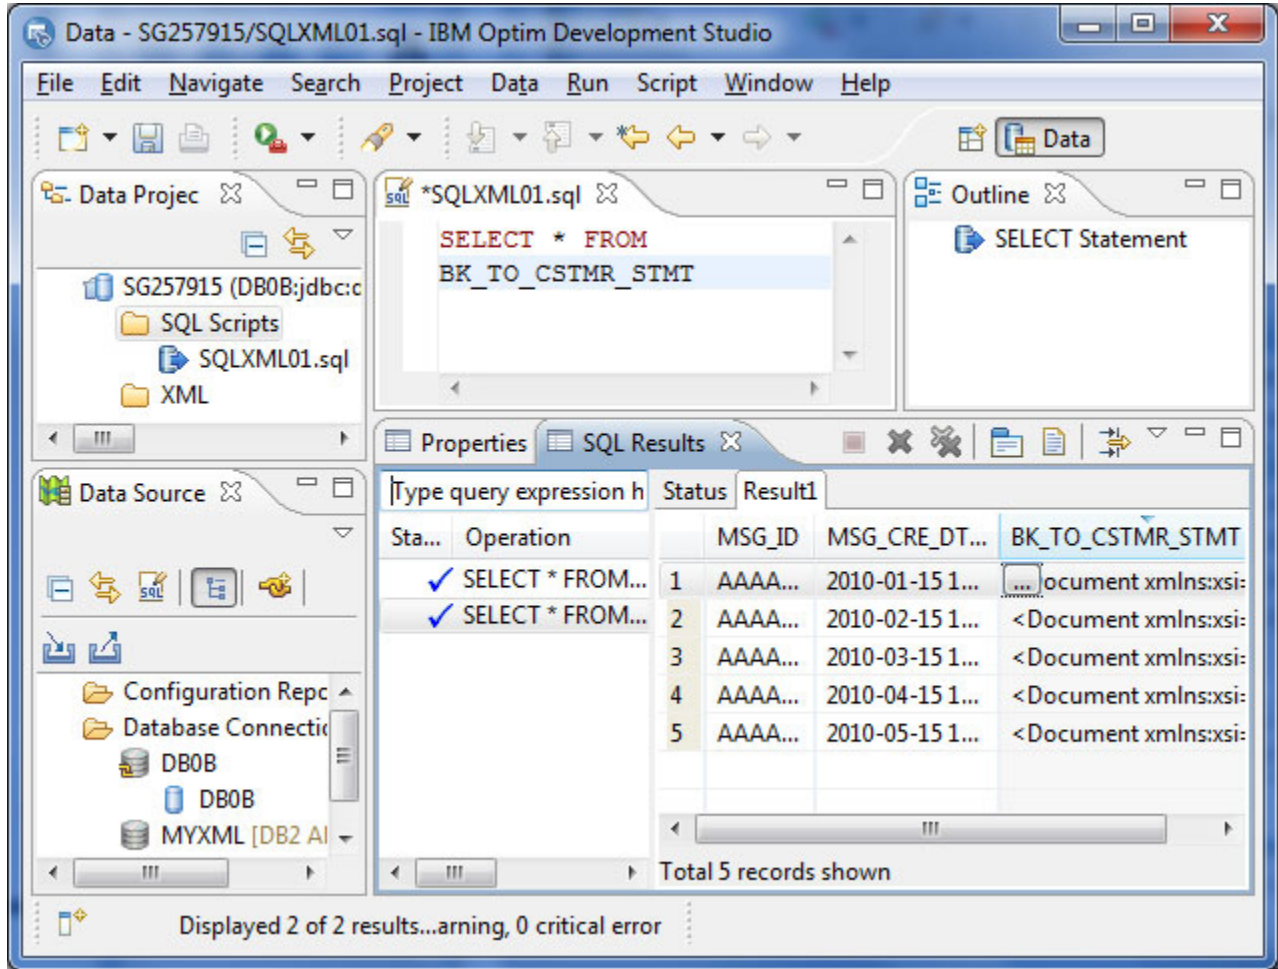


Figure 6-1 SQL Query - SELECT * FROM BK_TO_CSTMR_STMT

The query was executed through Optim Development Studio.

The contents of the various panes are:

- ▶ The top left pane is the project explorer, where we save our source code like SQL statements and stored procedure.
- ▶ The bottom left pane shows our database connection - to DB0B.
- ▶ The middle top pane shows the SQL statement, which we can “run” from the action bar
- ▶ The bottom right pane shows the results of the SQL statement.

The table contains the XML documents that we received from MQ, and the two fields that we stripped out using XMLQUERY and XMLTABLE functions in the stored procedure examples

If we wanted to view the hidden DOCID column too, we would need to explicitly select DB2_GENERATED_DOCID_FOR_XML column, along with MSG_ID, MSG_CRE_DT_TM and BK_TO_CSTMR_STMT columns.

Optim Development Studio also provides an XML document viewer. If we click on any of the BK_TO_CSTMR_STMT XML documents in the results pane, we can browse the contents of the XML document, and drill down to look at specific element and attribute values at any level within the XML document. The contents of the first XML document (MSG_ID AAAASESS-FP-STAT001) are illustrated in Figure 6-2.

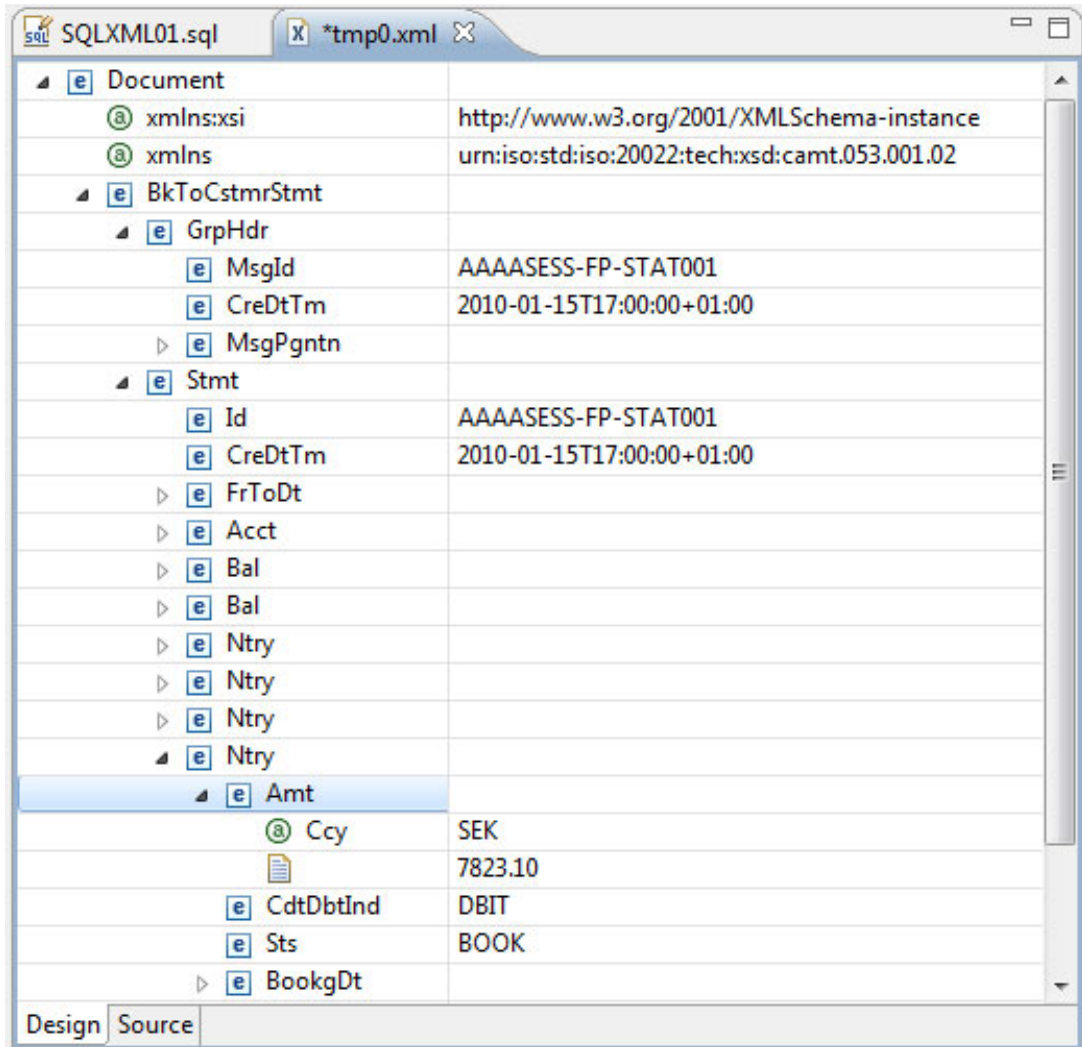


Figure 6-2 Optim Development Studio XML document viewer - Design view

The screen shot in Figure 6-2 shows that we have expanded the values of some data elements in the GrpHdr node. It also shows that the statement node this particular message has four Ntry elements, each representing a credit or debit transaction on the account, during the period covered by this statement. We could use this viewer to drill down and examine every single element in the XML document.

We could also click on the 'Source' tab at the bottom of the pane, and view 'Source' and 'Format' to ensure that the source view is formatted for easy viewing.

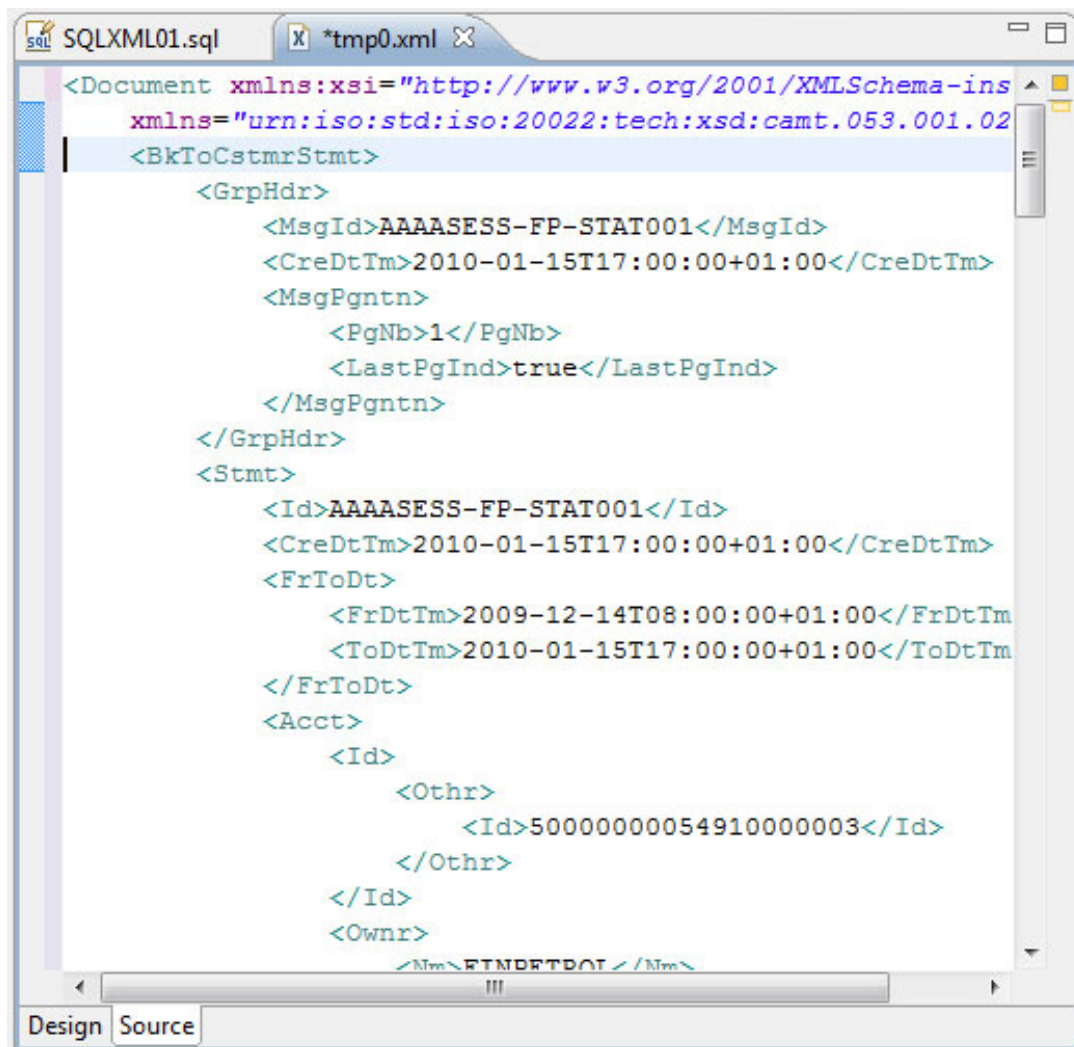


Figure 6-3 Optim Development Studio XML document viewer - Source view

Conceptually, the information contained in this XML document is exactly the same as what you see in your monthly current account statement from your own bank, but it looks very different in XML format. So, how easy would it be to use SQL/XML to query the Document and provide a more commonly recognizable view of the data?

The XMLTABLE function is designed for transforming elements from an XML document into a tabular view, similar to your monthly bank statement. The SQL/XML query in Example 6-13 retrieves all the debit and credit transactions from the XML document in a tabular format that looks much more like a traditional bank statement.

Example 6-13 SQL/XML Query to yield a "traditional" style bank statement

```
SELECT C.MSG_ID, C.MSG_CRE_DT_TM, X.BOOKED_DT_TM, X.AMT, X.CRDBTIND 1
FROM xmlr3.BK_TO_CSTMRT as C,
XMLTable(XMLNAMESPACES(
  DEFAULT 'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02'),2
  '$d/Document/BkToCstmrStmnt/Stmt/Ntry' PASSING C.BK_TO_CSTMRT as "d"
  COLUMNS 3
    "MSG_ID" VARCHAR(35) PATH '../GrpHdr/MsgId',
    "CRE_DT_TM" TIMESTAMP PATH '../GrpHdr/CreDtTm' ,
```

```

"BOOKED_DT_TM" TIMESTAMP PATH './BookgDt/DtTm' ,
"AMT" VARCHAR(50) PATH './Amt/text()' ,
"CRDDBTIND" VARCHAR(50) PATH './CdtDbtInd/text()' ) AS X
where C.MSG_ID = 'AAAASESS-FP-STAT001' 4
order by X.BOOKED_DT_TM asc ; 5

```

The highlighted numbered points of the SQL/XML query in Example 6-13 are explained as follows:

1. We are selecting three data elements from the XML document, which are the transaction timestamp, transaction amount and credit/debit indicator of each of the transactional entries in the BK_TO_CSTMTR_STMT XML document. Additionally, we are taking the statement message_id and statement date from the base table.
2. As always, we must specify the correct namespaces in the XQuery expression.
3. Important: we must specify the node within the XML document where the XMLTABLE function is operating on. Usually this will be at the node where we wish to extract the detailed information. In this case, we base it on the "Ntry" node, because we wish to retrieve a separate row in the result set for each transactional entry in the XML document.
4. Since we are focussing this query on a single XML document, we use a predicate on the base table column to return data from only the document with MSG_ID of AAAASESS-FP-STAT001.
5. Finally, we can order the result set using either a relational column, or an XML data element. In this case we wish to order the result set by the booked transaction date of each "Ntry" node in the XML document.

The results are displayed in Figure 6-4. Note that all rows returned contain the same values for columns from the base table (relating to the bank statement XML document) but unique values for the columns relating to individual entries from with the document.

Status	Result1	MSG_ID	MSG_CRE_DT_TM	BOOKED_DT_TM	AMT	CRDDBTIND
1		AAAASESS-FP-STAT001	2010-01-15 17:00:00.0	2010-01-18 12:15:00.0	23280	CRDT
2		AAAASESS-FP-STAT001	2010-01-15 17:00:00.0	2010-01-19 09:15:00.0	2023.44	DBIT
3		AAAASESS-FP-STAT001	2010-01-15 17:00:00.0	2010-01-23 09:15:00.0	833.99	DBIT
4		AAAASESS-FP-STAT001	2010-01-15 17:00:00.0	2010-01-27 09:15:00.0	7823.10	DBIT
Total 4 records shown						

Figure 6-4 Tabular result set of bank statement entries

The contents of XML data from multiple different documents can be returned in a single relational result set, as show by the SQL/XML statement in Example 6-14, which retrieves all 5 bank statements (January to May) and produces a consolidated bank statement for the entire period.

Example 6-14 SQL/XML query spanning multiple XML documents

```

SELECT C.MSG_ID, X.BOOKED_DT_TM, X.AMT, X.CRDDBTIND
FROM xmlr3.BK_TO_CSTMTR_STMT as C,
XMLTable(XMLNAMESPACES(
DEFAULT 'urn:iso:std:iso:2002:tech:xsd:camt.053.001.02'),

```

```
'$d/Document/BkToCstmrStmt/Stmt/Ntry' PASSING C.BK_TO_CSTMRT_STMT as "d"
COLUMNS
  "MSG_ID" VARCHAR(35) PATH '../GrpHdr/MsgId' ,
  "CRE_DT_TM" TIMESTAMP PATH '../GrpHdr/CreDtTm' ,
  "BOOKED_DT_TM" TIMESTAMP PATH './BookgDt/DtTm' ,
  "AMT" VARCHAR(50) PATH './Amt' ,
  "CRDBTIND" VARCHAR(50) PATH './CdtDbtInd'
) AS X
order by X.BOOKED_DT_TM asc ;
```

The results of the query in Example 6-14 are shown in Figure 6-5.

Status	Result1	MSG_ID	BOOKED_DT_TM	AMT	CRDBTIND
1		AAAASESS-FP-STAT001	2010-01-18 12:15:00.0	23280	CRDT
2		AAAASESS-FP-STAT001	2010-01-19 09:15:00.0	2023.44	DBIT
3		AAAASESS-FP-STAT001	2010-01-23 09:15:00.0	833.99	DBIT
4		AAAASESS-FP-STAT001	2010-01-27 09:15:00.0	7823.10	DBIT
5		AAAASESS-FP-STAT002	2010-02-18 12:15:00.0	23280	CRDT
6		AAAASESS-FP-STAT002	2010-02-19 09:15:00.0	2023.44	DBIT
7		AAAASESS-FP-STAT002	2010-02-23 09:15:00.0	384.30	DBIT
8		AAAASESS-FP-STAT002	2010-02-27 09:15:00.0	28982.80	DBIT
9		AAAASESS-FP-STAT002	2010-03-01 09:15:00.0	2928.00	DBIT
10		AAAASESS-FP-STAT002	2010-03-09 09:15:00.0	933.00	DBIT
11		AAAASESS-FP-STAT003	2010-03-18 12:15:00.0	23280	CRDT
12		AAAASESS-FP-STAT003	2010-03-19 09:15:00.0	2023.44	DBIT
13		AAAASESS-FP-STAT003	2010-03-23 09:15:00.0	722.50	DBIT
14		AAAASESS-FP-STAT003	2010-03-27 09:15:00.0	22929.90	DBIT
15		AAAASESS-FP-STAT004	2010-04-18 12:15:00.0	23280	CRDT
16		AAAASESS-FP-STAT004	2010-04-19 09:15:00.0	2023.44	DBIT
17		AAAASESS-FP-STAT004	2010-04-23 09:15:00.0	1398.00	DBIT
18		AAAASESS-FP-STAT004	2010-04-27 09:15:00.0	7338.77	DBIT
19		AAAASESS-FP-STAT004	2010-05-02 09:15:00.0	7290.00	DBIT

Total 24 records shown

Figure 6-5 Relational result set spanning data elements from multiple XML documents

A wider range of SQL/XML query techniques using all the DB2 XML functions will be illustrated in Section 7.4 SQL/XML Query techniques.

6.3.2 Choosing XML indexes

With the grand total of 5 rows in our table, indexes are not necessary. However, if we were storing a large number of XML documents in our XML auditing system, it would be essential to create XML indexes to support efficient searching within the XML documents.

As the XML audit database grows in size, the need for XML indexes will be clear. Each of the SQL/XML queries in 6.3.1, "Simple SQL/XML search examples" on page 99 will be

considered, and candidate XML indexes will be identified. Chapter 11, “Performance considerations” on page 243 covers XML index design in detail.

Let us consider XML index design for the specific purpose of enabling efficient search with an audit database of bank statements.

Let us assume that one of the primary purposes of the audit database will be to analyze data at the individual banking transaction level. That means that the focus of many queries will be the “Ntry” nodes in these XML documents. The typical data elements within an “Ntry” node are shown in Figure 6-6.

Element Name	Value
Ntry	
Amt	
Ccy	SEK
Text	7823.10
CdtDbtInd	DBIT
Sts	BOOK
BookgDt	
DtTm	2010-01-27T10:15:00+01:00
ValDt	
Dt	2010-01-27
AcctSvcrRef	AAAASESS-FP-ACCR-01
BkTxCd	
Domn	
Cd	PAYM
Fmly	
Cd	0001
SubFr	0003
NtryDtls	
Btch	
MsgId	FINP-0055
PmtInfId	FINP-0055/001
NbOfTxs	20

Figure 6-6 Typical “Ntry” node within a *Bk_To_Cstmr_Stmt* document

Let us assume that it was decided that many audit queries were to be focussed on the AccountServicerReference. We would check the ISO20022 documentation to understand the constraints on this data element, which are as follows:

- ▶ AccountServicerReference
- ▶ XPath = /Document/BkToCstmrStmt/Stmt/Ntry/AcctSvcrRef
- ▶ Presence: [0..1]
- ▶ Definition: Unique reference as assigned by the account servicing institution to unambiguously identify the entry.
- ▶ Data Type: Max35Text
- ▶ Format: maxLength: 35

- ▶ minLength: 1

There may be many occurrences of AcctSvcrRef in each Bk_To_Cstmr_Stmt document. An XML index on /Document/BkToCstmrStmt/Stmt/Ntry/AcctSvcrRef will therefore have multiple entries per Document. A candidate index definition is shown in Example 6-15.

Example 6-15 Candidate XML index definitions

```

Create index AcctSvcrRef_IX
  ON BK_TO_CSTMR_STMT(BK_TO_CSTMR_STMT)
  Generate Key using XMLPATTERN 'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
  /Document/BkToCstmrStmt/Stmt/Ntry/AcctSvcrRef'
as SQL VARCHAR(35) ;

```

6.3.3 Verifying XML index usage

Accurate data typing is critical for index eligibility. If the actual contents of the indexed data element in the document does not match the data type specified in the xmlpattern of the create index statement, then the index will not contain an entry for that data element.

The corollary of the previous statement is that it is very important for XML schemas to place data type constraints on fields that may be indexed, and the XML documents are validated against their schema definition. This is the best way of ensuring that effective indexes can be defined on XML columns.

The simplest way of checking that an Index is going to be effective is to run RUNSTATS and check the values of FIRSTKEYCARDF and FULLKEYCARDF in SYSIBM.SYSINDEXES. If you know the number of rows in the table, and you have a reasonable expectation of the average number of XML index hits per XML document, then you should have a rough idea of what the value of FULLKEYCARDF should be. For example, if you have 1 million rows in the BK_TO_CSTMR_STMT table, and an average of 20 “Ntry” nodes per XML document, then you might be expecting FULLKEYCARDF to be about 20 million. Of course it may be less if the same values of AcctSvcrRef crop up many times, but at least you can make a judgement on whether the Index has found approximately the right number of data elements to index.

If an XML index has cardinality of 0, then you should be asking yourself why the index did not find any matching data elements to index. The most likely reasons would be a typographical error in the xmlpattern, or the namespace definition, or a data type mismatch.

Once you have determined that the index has got approximately the correct number of entries in it, then you must test your application SQL to see whether the optimizer will use the XML index. This is a simple case of using your favorite explain tool, and checking whether or not the XML index(es) are being used in the access path for the query.

6.4 SQL/XML query techniques

Before considering more query techniques, let’s review the productivity that pureXML is providing.

The extent of the work that we have done so far is:

1. Download a publicly-available XML schema and register it in DB2.

2. Create a very simple DB2 table (with automatic schema validation) to store the corresponding XML documents that we receive.
3. "Load" the XML documents into the table (using the term "load" very loosely, to cover many possible ways to ingest the XML documents).
4. Start writing SQL/XML queries against a very rich XML data model
5. Create an iXML index to support those queries

We did not need to extract XML data elements from the XML documents to build relational searching columns. The only reason we chose to use the DB2 XML functions to extract selected data elements was to show how easy it was to move data between XML and relational.

The only difficult things we have had to do is to understand the concepts of the XML model of data, and the SQL/XML extensions to the SQL language.

This section provides some sample SQL/XML audit queries to illustrate the various functions and techniques for querying XML data and how to combine them.

6.4.1 Manipulating XML data with XPath functions

DB2 offers a wide range of functions for use in XPath expressions, these are a subset of the XPath 2.0 standard. They allow you to work directly on the data within the XPath expression instead of first having to extract the data from the document, and then manipulate or query it afterwards.

The functions include:

- ▶ Accessory functions, e.g. fn:data
- ▶ Functions on numeric values, e.g. fn:abs
- ▶ Functions on strings, e.g. fn:substring
- ▶ Functions on boolean values, e.g. fn:not
- ▶ Functions on durations, dates and times, e.g. fn:month-from-date
- ▶ Functions on sequences, e.g. fn:distinct-values
- ▶ Aggregate functions, e.g. fn:avg
- ▶ Context functions, e.g. fn:position

These functions all belong to the namespace <http://www.w3.org/2005/xpath-functions> and have the default namespace prefix of fn.

Additionally, constructor functions for each XPath data type are available. Examples of these are xs:string and xs:date. The namespace for these functions is <http://www.w3.org/2001/XMLSchema> and the prefix is xs.

For more information on the XPath functions available, refer to:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.xml1/db2z_xpxqfunctionreference.htm

Imagine that we want to calculate the sum of all the entries in a bank statement. This could be done by using the function fn:sum as shown in Example 6-16.

Example 6-16 Calculating the sum of the entries in a BankToCustomerStatement

```
SELECT
XMLCAST ( XMLQUERY (
'declare default element namespace
"urn:iso:std:iso:2002:tech:xsd:camt.053.001.02";
fn:sum(/Document/BkToCstmrStmnt/Stmnt/Ntry/Amt)')
```

```
PASSING BK_TO_CSTMRTMT
) AS DECIMAL(12,2)
FROM BK_TO_CSTMRTMT
```

The result of the XMLQUERY call is cast as DECIMAL(12,2) using the XMLCAST function. Otherwise the result of the query would have had type XML, despite the fact that we know that the content is really numeric. In case we want to process the data further using arithmetic functions, comparisons or just return it as a numeric value, we need to perform this conversion.

The use of the XMLCAST function requires that the input is a sequence of one item. If the result of the XMLQUERY expression is a longer sequence or an empty sequence, the XMLCAST function returns an error, SQL code -16003. In this example we know that the sequence has exactly one element, because the fn:sum function is applied to all the Amt elements so it can't be longer than one, and the result of applying fn:sum to an empty sequence is 0 so it always returns a result.

The XMLCAST function can also be used to strip any element tags if the result is a simple element.

As another example of applying XPath functions, let us consider some of the date functions available. Date and time functionality has been included in DB2 XPath with version 10 as this is an important aspect of most business applications.

Imagine that we wanted to select all the entries from a BankToCustomerStatement that were made in the last month. To create an XPath expression that does that, we shall use the subtraction function for datetime which is written as '-', and the two constructor functions xs:dateTime and xs:yearMonthDuration. The result is shown in Example 6-17.

Example 6-17 Using time and date functions in an XPath expression

```
SELECT
  XMLQUERY (
    'declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    Document/BkToCstmrtmt/Stmt/Ntry
    [BookgDt/DtTm>xs:dateTime($tm)-xs:yearMonthDuration("P1M")] '
    PASSING BK_TO_CSTMRTMT, CURRENT_TIMESTAMP AS "tm"
  )
FROM BK_TO_CSTMRTMT
```

The predicate says that the DtTm element must be greater than current timestamp minus one month, where the current timestamp is given as a parameter in the passing clause, and where both this parameter and the 1 month constant is created using constructor functions. The predicate is applied to each Ntry element in the document, and when evaluated to true, the Ntry element will be included in the result. This is because even though we reference the elements BookgDt and DtTm, these are only part of the predicate and not of the path that we have selected.

6.4.2 Filtering the rows returned with XMLEXISTS

The query shown in Example 6-17 returns one row per original table row, regardless of whether any entries in the BankToCustomerStatement qualified or not. For those BankToCustomerStatements that did not contain an entry less than a month old, the result of the query is the empty sequence. This is because predicates used in an XMLQUERY

expression are used to filter the data returned from the expression, not to determine whether data is returned at all.

In order to return only the rows that actually has some contents, we can apply the same filtering predicate in the where clause through the function XMLEXISTS. This function returns false for an empty sequence and true for everything else. It has the same syntax as XMLQUERY.

The resulting query is shown in Example 6-18.

Example 6-18 Avoiding empty sequences in result by using XMLEXISTS

```

SELECT
  XMLQUERY (
    'declare default element namespace
     "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
     Document/BkToCstmrStmt/Stmt/Ntry
     [BookgDt/DtTm>xs:dateTime($tm)-xs:yearMonthDuration("P1M")] '
    PASSING BK_TO_CSTMRT, CURRENT_TIMESTAMP AS "tm"
  )
FROM BK_TO_CSTMRT
WHERE XMLEXISTS (
  'declare default element namespace
   "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
   Document/BkToCstmrStmt/Stmt/Ntry
   [BookgDt/DtTm>xs:dateTime($tm)-xs:yearMonthDuration("P1M")] '
  PASSING BK_TO_CSTMRT, CURRENT_TIMESTAMP AS "tm"
)

```

XMLEXISTS predicates are candidates for index access if the XPath expression matches the pattern of an XML index.

Note: A predicate should always be in square brackets. If the brackets are omitted, the result of the XPath expression is either *true* or *false*.

Both true and false are non-empty sequences and would cause the XMLEXISTS expression to evaluate to true, thus returning all rows in the table.

6.4.3 Creating documents with publishing functions

In Example 6-18 the result is a sequence of Ntry elements for each resulting row rather than an XML document. If we want to create an XML document from the result, we can use the XML publishing functions XMLELEMENT and XMLDOCUMENT.

This is shown in Example 6-19 XMLELEMENT is used to create an outermost element with name Result, and XMLDOCUMENT to create a document from this element.

Example 6-19 Combining XMLQUERY with publishing functions

```

SELECT
XMLDOCUMENT( XMLELEMENT(NAME "results",
  XMLQUERY (
    'declare default element namespace
     "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
     Document/BkToCstmrStmt/Stmt/Ntry
     [BookgDt/DtTm>xs:dateTime($tm)-xs:yearMonthDuration("P1M")] '

```

```

    PASSING BK_TO_CSTMRT_STMT, CURRENT_TIMESTAMP AS "tm"
  )
))
FROM BK_TO_CSTMRT_STMT
WHERE XMLEXISTS (
  'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
  Document/BkToCstmrtStmt/Stmt/Ntry
  [BookgDt/DtTm>xs:dateTime($tm)-xs:yearMonthDuration("P1M")] '
  PASSING BK_TO_CSTMRT_STMT, CURRENT_TIMESTAMP AS "tm"
)

```

The result from this query is an XML document for each BankToCustomerStatement that had any entries the last month, containing these entries.

For more information on the use of publishing functions, refer to the following:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.xml/db2z_publishfuncs.htm
<http://www.ibm.com/developerworks/data/library/techarticle/dm-0511melnyk/>

6.4.4 Aggregating documents with XMLAGG

Now assume that instead of having individual XML documents for each BankToCustomerStatement with entries from the last month, we would like all these entries collected in one XML document regardless of their origin.

In this case we need to be able to create new XML elements *across* existing XML documents, and that is exactly what is offered by the publishing function XMLAGG. This function will take any number of XML values and create a sequence from them. It is an aggregate function in the same manner as the SQL functions AVG and MIN, so it is applied to all values from all rows in the select statement.

Example 6-20 shows how to combine Example 6-19 on page 109 with the XMLAGG function to collect all the entries in one XML document.

Example 6-20 Using XMLAGG to consolidate all entries into one document

```

SELECT
XMLDOCUMENT(
XMLAGG (
XMLELEMENT(NAME "results",
XMLQUERY (
'declare default element namespace
"urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
Document/BkToCstmrtStmt/Stmt/Ntry
[BookgDt/DtTm>xs:dateTime($tm)-xs:yearMonthDuration("P1M")] '
PASSING BK_TO_CSTMRT_STMT, CURRENT_TIMESTAMP AS "tm"
)
)))
FROM BK_TO_CSTMRT_STMT
WHERE XMLEXISTS (
'declare default element namespace
"urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
Document/BkToCstmrtStmt/Stmt/Ntry
[BookgDt/DtTm>xs:dateTime($tm)-xs:yearMonthDuration("P1M")] '

```

```
PASSING BK_TO_CSTMTR_STMT, CURRENT_TIMESTAMP AS "tm"
)
```

Note, that although we no longer get any rows containing empty sequences because all the Ntry elements are collected into one document, we still include the XMLEXISTS predicate to allow for index access.

6.4.5 Enumerating all occurrences using XMLTABLE

Now imagine that instead of creating one XML document with all the newest entries, we would instead like one XML document per entry. The function XMLTABLE can help us provide this.

The XMLTABLE function takes the following input:

- ▶ An optional namespace declaration
- ▶ A row XPath expression
- ▶ A number of column XPath expressions
- ▶ Optionally one or more input arguments as XPath variables

The row XPath expression returns a sequence of items, each of these produce a row in the result table of the XMLTABLE function. If the row XPath expression points to an element, the number of elements within a document with that particular XPath, determines the number of resulting rows. This is what allows us to create one row per Ntry element.

Using the row XPath expression as a starting point, one or more column XPath expressions define the contents of the columns returned from the XMLTABLE function. With each of these is associated a data type and a name which can be used in the surrounding select statement.

In Example 6-21 we show how to use XMLTABLE to extract all the entries no more than a month old, returning one row per entry. We use only one XPath column expression, namely '.' which effectively returns the contents of the row XPath expression. This is an Ntry element; it is returned with data type XML and subsequently made into an XML document using the publishing function XMLDOCUMENT.

Example 6-21 Extracting one entry per row using XMLTABLE

```
SELECT XMLDOCUMENT (X.NTRY)
FROM BK_TO_CSTMTR_STMT S,
XMLTABLE (XMLNAMESPACES(DEFAULT
    'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02',
    '/Document/BkToCstmtrStmt/Stmt/Ntry
[BookgDt/DtTm>xs:dateTime($tm)-xs:yearMonthDuration("P1M")]')
PASSING S.BK_TO_CSTMTR_STMT, current timestamp as "tm"
COLUMNS
"NTRY"      XML      PATH '.'
) X
```

Note, that we have now abandoned the XMLEXISTS predicate in the where clause. The row XPath expressions of an XMLTABLE function is a candidate for index access, so the XMLEXISTS predicate would yield exactly the same result and is no longer needed.

For more information and examples of the use of the XMLTABLE function, refer to

<http://www.ibm.com/developerworks/data/library/techarticle/dm-0708nicola/>

6.4.6 Grouping data with XMLTABLE

Finally, let us consider a case where we want to return all the entries ordered into documents according to the currency in which the entry was made. So we want one document with all the entries made in USD, one with all the entries made in SEK, and so on.

For each of the resulting documents we need entries from several of the original documents, but assuming that it is possible to make entries in different currencies to the same account, we cannot be sure that all entries from one account or indeed one BankToCustomerStatement will go into the same result document.

This can be obtained by combining the XMLTABLE function from Example 6-21 on page 111 with column grouping, XMLAGG and publishing functions XMLELEMENT and XMLDOCUMENT. The query is shown in Example 6-22.

We have added another column XPath expression to obtain the currency from the entry, and the result is grouped according to this currency. We then apply the XMLAGG function to the resulting entries, and wrap these in an element name Result. This in turn is then given as input to the XMLDOCUMENT function to produce an XML document.

Example 6-22 Grouping entries obtained from XMLTABLE according to currency

```

SELECT XMLDOCUMENT (
  XMLELEMENT(NAME "Result",
    XMLAGG(X.NTRY))
  FROM BK_TO_CSTMRTMT S,
  XMLTABLE (XMLNAMESPACES(DEFAULT
    'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02',
    '/Document/BkToCstmrtmt/Stmt/Ntry
    [BookgDt/DtTm>xs:dateTime($tm)-xs:yearMonthDuration("P1M")]')
  PASSING S.BK_TO_CSTMRTMT, current timestamp as "tm"
  COLUMNS
    "CCY"          VARCHAR(20)  PATH 'Amt/@Ccy' ,
    "NTRY"        XML          PATH '.'
  ) X
  GROUP BY X.CCY

```

6.5 User defined functions with XML

User defined functions are a great way of encapsulating complex logic or expressions in function that is very easy for anybody to use. The first stored procedure in this chapter (which received an XML document, extracted a number of relational and XML objects from the received XML document, and stored the results in a DB2 table) contained some XQuery expressions which would be quite challenging for an SQL programmer to use if that programmer had not been exposed to XML.

6.5.1 UDFs for reading from XML documents

We could dramatically simplify that stored procedure by providing three user defined functions to perform the XML tasks. In Example 6-23 we provide the DDL necessary to create the three UDFs that would really help the traditional SQL programmer.

Example 6-23 Creating three user defined functions

```

CREATE FUNCTION GETMSGID(doc XML) RETURNS VARCHAR(35)
  LANGUAGE SQL
  CONTAINS SQL
  NO EXTERNAL ACTION
  DETERMINISTIC
  DISABLE DEBUG MODE
BEGIN
RETURN xmlcast(xmlquery(
  'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
  $d/Document/BkToCstmrStmt/GrpHdr/MsgId' passing doc as "d")
  as varchar(35) ) ;
END !

CREATE FUNCTION GETCREDITM(doc XML) RETURNS TIMESTAMP
  LANGUAGE SQL
  CONTAINS SQL
  NO EXTERNAL ACTION
  DETERMINISTIC
  DISABLE DEBUG MODE
BEGIN
RETURN xmlcast(xmlquery(
  'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
  $d/Document/BkToCstmrStmt/GrpHdr/CreDtTm' passing doc as "d")
  as timestamp ) ;
END !

CREATE FUNCTION GETMINISTMT(doc XML) RETURNS XML
  LANGUAGE SQL
  CONTAINS SQL
  NO EXTERNAL ACTION
  DETERMINISTIC
  DISABLE DEBUG MODE
BEGIN
RETURN xmlquery(
  'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
  $d/Document/BkToCstmrStmt/Stmt' passing doc as "d") ;
END !

```

Each of these three UDFs encapsulates XQuery functions to return a different data type to the SQL programmer. Example 6-24 shows how these functions can be used in standard SQL.

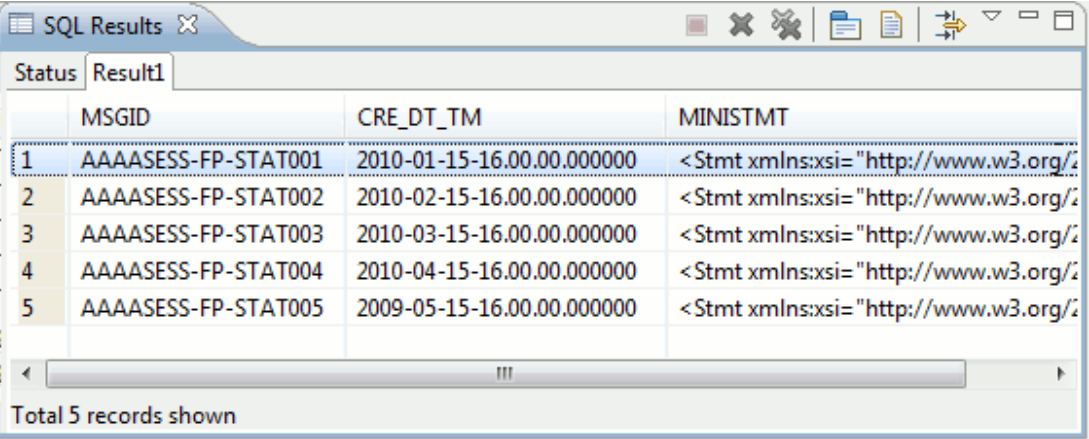
Example 6-24 Usage of UDFs

```

select
  getmsgid(BK_TO_CSTMRTMT) as MSGID,
  getcredttm(BK_TO_CSTMRTMT) as CRE_DT_TM,
  getministmt(BK_TO_CSTMRTMT) as MINISTMT
from BK_TO_CSTMRTMT

```

The results of the UDF-based query in Example 6-24 are shown in Figure 6-7.



Status	Result1		
	MSGID	CRE_DT_TM	MINISTMT
1	AAAASESS-FP-STAT001	2010-01-15-16.00.00.000000	<Stmt xmlns:xsi="http://www.w3.org/2003/XMLSchema-instance" type="statement" id="1"
2	AAAASESS-FP-STAT002	2010-02-15-16.00.00.000000	<Stmt xmlns:xsi="http://www.w3.org/2003/XMLSchema-instance" type="statement" id="2"
3	AAAASESS-FP-STAT003	2010-03-15-16.00.00.000000	<Stmt xmlns:xsi="http://www.w3.org/2003/XMLSchema-instance" type="statement" id="3"
4	AAAASESS-FP-STAT004	2010-04-15-16.00.00.000000	<Stmt xmlns:xsi="http://www.w3.org/2003/XMLSchema-instance" type="statement" id="4"
5	AAAASESS-FP-STAT005	2009-05-15-16.00.00.000000	<Stmt xmlns:xsi="http://www.w3.org/2003/XMLSchema-instance" type="statement" id="5"

Total 5 records shown

Figure 6-7 SQL Results using UDFs on XML documents

So, if we rewrite the stored first stored procedure to use these functions (as shown in Example 6-25) we find that there is it now looks much simpler. In fact, the traditional SQL programmer does not need to know a shred of XQuery, provided that they have a range of system provided and user defined functions to operate on XML data.

Example 6-25 Modified stored procedure using UDFs instead of XQuery expressions

```

CREATE PROCEDURE STOREXML5 (
    IN V_BANKSTMT XML,
    OUT V_MSG_ID VARCHAR(35),
    OUT V_CREDITM TIMESTAMP,
    OUT V_MINISTMT XML)
    LANGUAGE SQL
    MODIFIES SQL DATA
    DISABLE DEBUG MODE

BEGIN

DECLARE VALIDXML XML ;
DECLARE SQLCODE INTEGER ;

SET VALIDXML = (
    SELECT DSN_XMLVALIDATE(V_BANKSTMT, 'SYSXSR.CAMT_053_001_02');

SELECT
    getmsgid(BK_TO_CSTMRT_STMT),
    getcredttm(BK_TO_CSTMRT_STMT),
    getministmt(BK_TO_CSTMRT_STMT)
    into V_MSG_ID, V_CREDITM, V_MINISTMT from BK_TO_CSTMRT_STMT;

INSERT INTO BK_TO_CSTMRT_STMT_MANUALVALIDATE (
    MSG_ID , MSG_CRE_DT_TM, BK_TO_CSTMRT_STMT)
    values ( V_MSG_ID, V_CREDITM, VALIDXML ) ;

END !

```

UDFs for XML are a great way to take advantage of DB2's pureXML capabilities, without necessarily exposing the SQL programmers to the full complexity of XQuery expressions.

6.5.2 UDFs for writing updates to XML documents

The UDF examples we have shown so far have been based around retrieving data from XML documents, and transforming them to traditional relational objects. UDFs can also be used for other purposes, such as writing sub-document updates.

Example 6-26 illustrates a user defined function that inserts a new XML node into an XML document.

Example 6-26 UDF for XML sub-document update

```
CREATE FUNCTION UPDATE_ADDR(NEW_ADDR CLOB)
  RETURNS integer
  LANGUAGE SQL
  MODIFIES SQL DATA
  DETERMINISTIC
  DISABLE DEBUG MODE
BEGIN
UPDATE CUSTOMER_TABLE
  SET CUSTOMER_ADDRESS = XMLMODIFY(
    'insert node $new as last into /Customer/addresses',
    XMLPARSE(NEW_ADDR) as "new" );
  return SQLCODE;
END !

-- This UDF could be called using the following SQL statement -

SELECT UPDATE_ADDR(V_ADDR) INTO UDF_RETURN1
  FROM CUSTOMER_TABLE
  where CUSTID = 'VALUE' ;
```

6.6 Triggers with XML

Triggers are a popular way of implementing database dependencies. Triggers can be defined to perform procedural work as a direct result of an SQL write (insert, update or delete) to the table that the trigger is defined on. Triggers can use "transition variables" to reference the new or old values of DB2 columns, which allows the trigger routine body to work with the data values that have just been written.

Triggers can be used on tables with XML columns. However, the new and old values of XML columns are not available as transition variables to the trigger routine body.

In summary, you can continue to write triggers with tables that have XML columns, but if you want to reference the contents of an XML column you will need to re-read the XML data using one of the relational columns that are available as a transition variable.

6.7 XML joins

We examine XML to relational join and XML to XML join.

6.7.1 XML to relational join

Joining XML data with relational data is no different in principle to wholly relational joins. The data types of the join must be compatible, and indexes should be used to make the join perform well.

The ISO20022 sample XML messages that we are using all relate to an account called 'FINPETROL'. The account name is stored at the XPath location:

```
/Document/BkToCstmrStmnt/Stmnt/Acct/Ownr/Nm
```

We can perform a join based on the Account Name within the XML document, against a relational table of addresses, illustrated in Example 6-27.

Example 6-27 Contents of relational address table

CUSTNAME	STRTNM	BLDGNB	PSTCD	TWNNM
FINPETROL	Bailey Avenue	555	95141	San Jose
WINGPETROL	Bond Street	23	98282	New York
TAILPETROL	Southfork	1	99999	Dallas

There are many ways to perform an XML to relational join in SQL/XML. One approach is to use the XMLEXISTS function (which is indexable). A very simple join statement is illustrated in Example 6-28. The join is achieved by passing the relational column for the join (a.custname) to the XMLEXISTS function as a variable which is then used as an XML predicate.

Example 6-28 XML to Relational Join example using XMLEXISTS

```
select c.msg_id, c.msg_cre_dt_tm, a.strtnm, a.bldgnb, a.pstcd, a.twnnm
from BK_TO_CSTMRT_STMT c, ADDRESS a
  where xmlexists('declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $i/Document/BkToCstmrStmnt/Stmnt/Acct/Ownr[Nm=$acctname]'
    passing
      c.BK_TO_CSTMRT_STMT as "i",
      a.custname as "acctname");

--yields
```

MSG_ID	MSG_CRE_DT_TM	STRTNM	BLDGNB	PSTCD	TWNNM
AAAASESS-FP-STAT001	2010-01-15-17.00.00...	Bailey Avenue	555	95141	San Jose

Another method of performing an XML to relational join is using the XMLTABLE function, which is indexable if coded with a predicate, as shown in figure. Example 6-29 shows normal relational join being coded between the result of an XMLTABLE function and a relational table. In Example 6-29 there is no predicate on the XMLTABLE function, so this particular query would not be indexable.

Example 6-29 XML to relational join example using XMLTABLE

```

SELECT X.ACCT_NM, X.SVCR_NM, a.strtnm, a.bldgnb, a.pstcd, a.twnnm
FROM
  BK_TO_CSTMR_STMT as C,
  ADDRESS as a,
  XMLTable(XMLNAMESPACES(
    DEFAULT'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02'),
    '$d/Document/BkToCstmrStmt/Stmt/Acct'
    PASSING c.BK_TO_CSTMR_STMT as "d"
    COLUMNS "ACCT_NM"  VARCHAR(35)  PATH './Ownr/Nm',
    "SVCR_NM"  VARCHAR(35)  PATH './Svcr/FinInstnId/Nm' ) AS X
where x.ACCT_NM = a.CUSTNAME ;

-- yields

```

ACCT_NM	SVCR_NM	STRTNM	BLDGNB	PSTCD	TWNNM
FINPETROL	AAAA	BANKEN	Bailey Avenue	555	95141 San Jose

6.7.2 XML to XML join

The techniques required to perform XML to XML join are a little different. In order to provide a test case we will convert the relational address table (used in the XML to relational join examples in 6.7.1, "XML to relational join" on page 116) to an XML format. This allows us to perform the same logical joins, except that both sources are XML.

The relational address table is converted to XML using the script in Example 6-30.

Example 6-30 Script to convert the relational address table to XML

```

create table xmladdress ( CUSTADDR XML ) ;

INSERT INTO XMLADDRESS (CUSTADDR)
  SELECT XMLSERIALIZE(XMLDOCUMENT(
    XMLELEMENT(NAME "MsgRcpt",
    XMLELEMENT(NAME "Nm", CUSTNAME),
    XMLELEMENT(NAME "Pst1Adr",
      XMLELEMENT(NAME "StrtNm", STRTNM),
      XMLELEMENT(NAME "BldgNb", BLDGNB),
      XMLELEMENT(NAME "PstCd", PSTCD),
      XMLELEMENT(NAME "TwnNm", TWNNM) ))) AS CLOB)
from Address ;

select * from xmladdress ;

-- yields

CUSTADDR
-----
<MsgRcpt><Nm>FINPETROL</Nm><Pst1Adr><StrtNm>Bailey
Avenue</StrtNm><BldgNb>555</BldgNb><PstCd>95141</PstCd><TwnNm>San
Jose</TwnNm></Pst1Adr></MsgRcpt>

```

```
<MsgRcpt><Nm>WINGPETROL</Nm><Pst1Adr><StrtNm>Bond
Street</StrtNm><BldgNb>23</BldgNb><PstCd>98282</PstCd><TwnNm>New
York</TwnNm></Pst1Adr></MsgRcpt>
```

```
<MsgRcpt><Nm>TAILPETROL</Nm><Pst1Adr><StrtNm>Southfork</StrtNm><BldgNb>1</BldgNb><
PstCd>99999</PstCd><TwnNm>Dallas</TwnNm></Pst1Adr></MsgRcpt>
```

Having converted the address table to XML, we can now code XML to XML joins. Example 6-31 shows the use of an XMLEXISTS function to join the two tables. In this example we need to pass both XML documents (“i” and “j”) into the XMLEXISTS function to perform the comparison.

The XML string function is used with the current location to signify that we wish to perform a string comparison to join the current contents of one location in one document with the current contents of another location in the other document. The specification of the data type being used in the join comparison is very important because

- ▶ Unlike relational columns, the XML fields do not have types defined
- ▶ Index eligibility is dependent on the data type

Example 6-31 XML to XML join using XMLEXISTS

```
select * from BK_TO_CSTMR_STMT c, XMLADDRESS a
  where xmlexists('declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $i/Document/BkToCstmrStmnt/Stmnt/Acct/Ownr[
      Nm/fn:string(.) = $j/MsgRcpt[ Nm/fn:string(.) ] ]'
    passing
      c.BK_TO_CSTMR_STMT as "i",
      a.CUSTADDR as "j");
```

Example 6-31 was coded in order to illustrate the mechanics of an XML to XML join as clearly as possible. This query does not contain any filter predicate that is indexable. However, the join predicate is eligible for index access because it has been expressed with the string() function, which is indexable.

Similarly, if you have XML to XML join on a numeric field, you should use xs:double so that an XML DECFLOAT index can be used for the join predicate.

6.8 XML with change data capture tools

The primary application scenario for this book is the ISO20022 standard for banking messages.

Another common source of XML messages is the range of replication and event publishing tools that are used to capture changes from existing database, and publish change data capture (CDC) messages. The published messages would sometimes be used to replicate changes to another database (such as a data warehouse). It is becoming increasingly common that *CDC messages* (change data capture messages) are being used in event driven systems. Changes to source data, that meets certain criteria (say, bank transfers that exceed a threshold value) could be routed to a workflow system (such as WebSphere Message Broker) where the CDC event would be examined using workflows that implement business processes, and initiate automated actions.

This section considers some ways in which CDC messages could be handled by DB2 pureXML.

6.8.1 Change data capture tools background

Change data capture tools tend to follow an architecture along the following lines:

1. A *capture* process is used to read the database log of a source database, looking for changes which have been requested by a subscription definition.
2. When qualifying changes are found in the log, they are packaged up (typically into unit of work boundaries) and transmitted over a network infrastructure to the target systems that have subscribed to them.
3. The target systems receive the changes and do something with them (such as update a database, invoke an application process, publish an CDC message in XML or some other format, and so on).

IBM's replication and event publishing tools that publish XML change data capture messages include the tools listed in Table 6-2.

Table 6-2 IBM change data capture tools that publish XML messages.

Tool	Main data sources supported for CDC	Comments
InfoSphere Data Event Publisher	<ul style="list-style-type: none"> ▶ DB2 for z/OS ▶ DB2 for LUW ▶ Oracle 	<ul style="list-style-type: none"> ▶ Asynchronous log reader services ▶ Writes CDC messages as XML (and other formats) ▶ Publishes messages directly to WebSphere MQ ▶ z/OS and LUW versions
InfoSphere Classic Data Event Publisher	<ul style="list-style-type: none"> ▶ IMS ▶ VSAM ▶ IDMS ▶ Adabas 	<ul style="list-style-type: none"> ▶ Asynchronous log reader services ▶ Writes CDC messages as XML (and other formats) ▶ Publishes messages directly to WebSphere MQ or zFS files ▶ z/OS only
InfoSphere Change Data Capture	<ul style="list-style-type: none"> ▶ DB2 for z/OS ▶ DB2 for LUW ▶ DB2 for iSeries® ▶ Oracle ▶ Sybase ▶ SQL Server 	<ul style="list-style-type: none"> ▶ Asynchronous log reader services ▶ Writes CDC messages as XML (and other formats) ▶ Source server writes CDC data over tcpip to a target server. ▶ Target server writes messages to a range of targets, including WebSphere MQ and files ▶ z/OS and LUW versions

The purpose of this section is to focus on the XML messages that are published from these tools, and examine how they can be used with DB2 pureXML.

InfoSphere Data Event Publisher and InfoSphere Classic Data Event Publisher both share a common schema and generate messages like the one shown in Example 6-32.

Example 6-32 Sample XML CDC message format for DB2 and Classic Data Event Publishers

```
<?xml version="1.0" ?>
  <msg xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:noNamespaceSchemaLocation="mqcap.xsd"
    version="1.0.0" dbName="$IMS"
    seqNum="IMS000003B4009FF00AIMSY" _8"> 1
    <rowOp authID="USER0000"
      planName="USER0000"
      cmitLSN="IMS000003B4009FF00AIMSY"
      cmitTime="2010-11-12:14:22:02.374839"> 2
    <insertRow subName="CLASSIC" srcOwner="IMSP" srcName="CLASSIC1"> 3
```

```

    <col name="STRTNM" isKey="0"> 4
      <char>Bailey Avenue</char>
    </col>
    <col name="BLDGNB" isKey="0">
      <char>555</char>
    </col>
    <col name="PSTCD" isKey="0">
      <char>95149</char>
    </col>
    <col name="TWNNM" isKey="0">
      <char>San Jose</char>
    </col>
  </insertRow>
</rowOp>
</msg>

```

The XML message is quite easy to understand by reviewing it. The annotated points in the figure are:

1. The XML schema is mqcap.xsd is available to be defined in the pureXML XSR. the document root also contains information about the database source.
2. Change data capture messages can be generated based either as row operations (like this one) or as transactions. A transaction message inserts <trans> tags around one or more <rowOp> tags. log sequence numbers and timestamps are included as attributes, and reflect the information that is available for a particular data source.
3. Individual row operations may be <InsertRow>, <UpdateRow> or <DeleteRow>
4. The before and after images of the column values are published. The administration tools has options to ignore before images, and make other customization to the published data.

InfoSphere CDC offers the choice of a range of XML schemas for publishing change data messages. Example 6-33 is an SQL update example of a simple InfoSphere CDC XML message.

Example 6-33 Sample XML CDC message format for InfoSphere CDC

```

<?xml version="1.0" encoding="UTF-8"?>
  <CUST_CHANGE>
    <TimeStamp AfterImage="2010-10-22T14.22.38.000"
      BeforeImage="2009-05-16T10.09.22.000"/>
    <SourceLogSeqNumber AfterImage="20101022142238000000"
      BeforeImage="20090516100922000000"/>
    <STRTNM AfterImage=" Nice Street " BeforeImage=" Bailey Avenue "/>
    <BLDGNB AfterImage=" 100 " BeforeImage=" 555 "/>
    <PSTCD AfterImage=" 98282 " BeforeImage=" 95141 "/>
    <TWNNM AfterImage=" San Diego " BeforeImage=" San Jose "/>
  </CUST_CHANGE>

```

Based on the XML examples we have worked on so far in this chapter, these CDC XML messages look fairly straightforward to handle.

6.8.2 Using DB2 pureXML to receive CDC messages

Change data capture event messages are useful for many different reasons, such as

- ▶ updating a data warehouse with a stream of change data messages, so that it can be kept up to date with very low latency compared to other techniques such as extracts and loads.
- ▶ notifying workflow or event driven systems about important changes to operational data, that require rapid and/or automated business processes to be executed

Another potential use for CDC messages is to maintain a historical record of changes to operational data. We expand this idea in 6.8.3, “XML history objects” on page 125.

For now, let us examine how we would consume a stream of XML CDC messages with DB2 pureXML. We could use a mixture of all the programming constructs that we have illustrated in this chapter so far to receive CDC XML messages into DB2 pureXML storage, as illustrated in Figure 6-8.

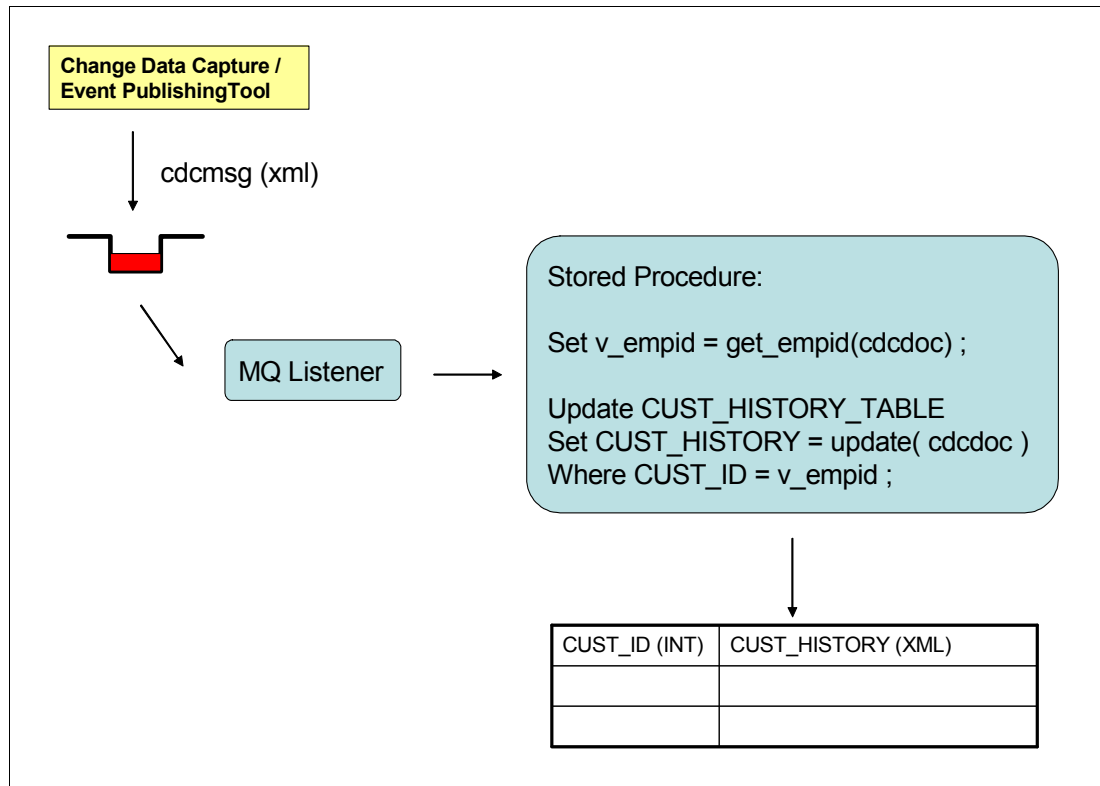


Figure 6-8 Scenario to receive XML CDC messages into DB2 pureXML via MQ

We have already discussed the mechanics of receiving XML messages from WebSphere MQ earlier in this chapter. The requirements will be more complex if we wish to update a historical audit record.

Figure 6-9 shows the contents of the XML historical record before the CDC record is applied. Customer John Doe, with customer number ‘CUST1’ has a record of two nominated email addresses (one expired, one active) and only one postal address (which is active). If we receive a CDC message containing a change of address then we will want to set an end date in the node for the current postal address, and add a new node for the new address.

```

<?xml version="1.0" encoding="UTF-8"?>
<CustomerHistory>
  <DB2CustomerDB>
    <customer_identification>
      <customer_id>CUST1</customer_id>
      <customer_name>John Doe</customer_name>
      <emails>
        <email>
          <email_effective_from>2008-07-18T00:00:00+01:00</email_effective_from>
          <email_effective_to>2009-10-18T00:00:00+01:00</email_effective_to>
          <email_addr>john_doe@isp.com</email_addr>
        </email>
        <email>
          <email_effective_from>2008-07-18T00:00:00+01:00</email_effective_from>
          <email_effective_to> </email_effective_to>
          <email_addr>jon@newisp.com</email_addr>
        </email>
      </emails>
      <addresses>
        <address>
          <address_effective_from>2008-07-18T00:00:00+01:00</address_effective_from>
          <address_effective_to> </address_effective_to>
          <Pst1Adr>
            <StrtNm> Bailey Avenue </StrtNm>
            <BldgNb> 555 </BldgNb>
            <PstCd> 95141 </PstCd>
            <TwnNm> San Jose </TwnNm>
          </Pst1Adr>
        </address>
      </addresses>
    </customer_identification>
  </DB2CustomerDB>
</CustomerHistory>

```

Figure 6-9 Initial CUST_HISTORY table contents for 'CUST1'

If we receive the CDC message in Example 6-33 we need to write a procedure to do the following:

1. Examine the contents of the XML CDC record to find the customer identifier key.
2. Use the customer identifier key to retrieve the XML document that stores the historical record of customer address changes.
3. Update the XML node in the document that stored the current address details, to set an end data for that address.
4. Insert a new node in the XML document to reflect the new address details

The stored procedure in Example 6-34 receives the XML CDC message from MQ (in the format of the InfoSphere CDC example), and updates the historical record of changes in DB2. This stored procedure actually picks up the CDC message from a test table, but the MQ examples earlier in the chapter show you how you would adapt it to work with WebSphere MQ.

Example 6-34 Stored Procedure to receive and apply CDC message

```

CREATE PROCEDURE XMLR3.RECEIVE_CDC(IN CDCDOC XML) 1
LANGUAGE SQL

```

```

MODIFIES SQL DATA
DISABLE DEBUG MODE

BEGIN

DECLARE CDC_CUSTID CHAR(5) ;
DECLARE CDC_STRTNM VARCHAR(70) ;
DECLARE CDC_BLDGNB VARCHAR(16) ;
DECLARE CDC_PSTCD VARCHAR(16) ;
DECLARE CDC_TWNNM VARCHAR(35) ;
DECLARE V_ADDR_FROM TIMESTAMP WITH TIMEZONE ;
DECLARE V_STRTNM VARCHAR(70) ;
DECLARE V_BLDGNB VARCHAR(16) ;
DECLARE V_PSTCD VARCHAR(16) ;
DECLARE V_TWNNM VARCHAR(35) ;
DECLARE PREV_MAX_ADDR XML ;
DECLARE NEXT_MAX_ADDR XML ;
DECLARE UDF_RETURN1 INT ;
DECLARE UDF_RETURN2 INT ;

SELECT X.CUSTID, X.STRTNM, X.BLDGNB, X.PSTCD, X.TWNNM
      INTO CDC_CUSTID, CDC_STRTNM, CDC_BLDGNB, CDC_PSTCD, CDC_TWNNM
      FROM XMLTable('$d/CUST_CHANGE' PASSING CDCDOC as "d"
                   COLUMNS
                     "CUSTID" CHAR(5) PATH 'CUSTID/@AfterImage',
                     "STRTNM" VARCHAR(70) PATH 'STRTNM/@AfterImage',
                     "BLDGNB" VARCHAR(16) PATH 'BLDGNB/@AfterImage',
                     "PSTCD" VARCHAR(16) PATH 'PSTCD/@AfterImage',
                     "TWNNM" VARCHAR(35) PATH 'TWNNM/@AfterImage'
                   ) AS X ; ❷

SELECT X.ADDR_FROM, X.STRTNM, X.BLDGNB, X.PSTCD, X.TWNNM
      INTO V_ADDR_FROM, V_STRTNM, V_BLDGNB, V_PSTCD, V_TWNNM
      FROM CUST_HISTORY C,
      XMLTable(
        '$cu/CustomerHistory/DB2CustomerDB/customer_identification/addresses/address'
        PASSING C.CUST_HISTORY_OBJECT as "cu"
        COLUMNS
          "ADDR_FROM" TIMESTAMP WITH TIMEZONE PATH 'address_effective_from',
          "STRTNM" VARCHAR(50) PATH 'Pst1Adr/StrtNm',
          "BLDGNB" VARCHAR(50) PATH 'Pst1Adr/BlDgNb',
          "PSTCD" VARCHAR(50) PATH 'Pst1Adr/PstCd',
          "TWNNM" VARCHAR(50) PATH 'Pst1Adr/TwnNm'
        ) AS X
      WHERE CUSTID = CDC_CUSTID
      order by X.ADDR_FROM desc
      fetch first row only ; ❸

SET PREV_MAX_ADDR = XMLELEMENT(NAME "address",
  XMLELEMENT(NAME "address_effective_from", V_ADDR_FROM),
  XMLELEMENT(NAME "address_effective_to", current timestamp with timezone),
  XMLELEMENT(NAME "Pst1Adr",
    XMLELEMENT(NAME "StrtNm", V_STRTNM),
    XMLELEMENT(NAME "BlDgNb", V_BLDGNB),
    XMLELEMENT(NAME "PstCd", V_PSTCD),

```

```

XMLLEMENT(NAME "TwnNm", V_TWNNM))) ; 4
SET NEXT_MAX_ADDR = XMLLEMENT(NAME "address",
XMLLEMENT(NAME "address_effective_from", current timestamp with timezone),
XMLLEMENT(NAME "address_effective_to", ''),
XMLLEMENT(NAME "Pst1Adr",
XMLLEMENT(NAME "StrtNm", CDC_STRTNM),
XMLLEMENT(NAME "BldgNb", CDC_BLDGNB),
XMLLEMENT(NAME "PstCd", CDC_PSTCD),
XMLLEMENT(NAME "TwnNm", CDC_TWNNM))) ; 5

UPDATE CUST_HISTORY X
SET X.CUST_HISTORY_OBJECT = XMLMODIFY(
'replace node
/CustomersHistory/DB2CustomerDB/customer_identification/addresses/address[
address_effective_from=$t]
with $V',
PREV_MAX_ADDR AS "V",
V_ADDR_FROM as "t" )
where CUSTID = 'CUST1' ; 6

UPDATE CUST_HISTORY X
SET X.CUST_HISTORY_OBJECT = XMLMODIFY(
'insert node $new as last into
/CustomersHistory/DB2CustomerDB/customer_identification/addresses',
NEXT_MAX_ADDR as "new" )
where CUSTID = 'CUST1' ; 7

END !

```

The logic of the stored procedure is explained as follows with the annotated points from Example 6-34.

1. The stored procedure accepts the cdc message as an XML document in INPUT parameter CDCDOC.
2. The data elements of the incoming cdc message are stripped with a single XMLTABLE function
3. Using the customer id ('CUST1') from the cdc document, the current address node is retrieved from the CUSTOMER_HISTORY table.
4. A replacement XML node is derived for the current address node, using XML Publishing functions
5. A new XML node is derived from the incoming CDCDOC
6. The existing XML node storing the current address is replaced using an XMLMODIFY function.
7. The new XML node is inserted using a further XMLMODIFY function.

PREV_MAX_ADDR and NEXT_MAX_ADDR are kept as XML type as they are going to be used in XMLMODIFY later, therefore avoiding an XMLPARSE().

The resultant contents of the addresses nodes in the customer history table is shown in Figure 6-10.

```

<addresses>
  <address>
    <address_effective_from>2008-07-18T00:00:00.000000+01:00</address_effective_from>
    <address_effective_to>2010-11-08T17:31:03.602776-05:00</address_effective_to>
    <PstlAdr>
      <StrtNm> Bailey Avenue </StrtNm>
      <BldgNb> 555 </BldgNb>
      <PstCd> 95141 </PstCd>
      <TwnNm> San Jose </TwnNm>
    </PstlAdr>
  </address>
  <address>
    <address_effective_from>2010-11-08T17:31:03.602935-05:00</address_effective_from>
    <address_effective_to/>
    <PstlAdr>
      <StrtNm> Nice Street </StrtNm>
      <BldgNb> 100 </BldgNb>
      <PstCd> 98282 </PstCd>
      <TwnNm> San Diego </TwnNm>
    </PstlAdr>
  </address>
</addresses>

```

Figure 6-10 Updated CUST_HISTORY table contents for 'CUST1'

If you use the InfoSphere Data Event Publisher or Classic Data Event Publisher products, the XML schema is different, but the technique is the same. All you need to do is replace the XMLTABLE operation on the incoming XML CDC document with a different XMLTABLE operation shown in Example 6-35, to work with the incoming schema.

Example 6-35 XMLTABLE function for Event Publisher XML schemas

```

SELECT X.CUSTID, X.STRTNM, X.BLDGNB, X.PSTCD, X.TWNNM
  INTO CDC_CUSTID, CDC_STRTNM, CDC_BLDGNB, CDC_PSTCD, CDC_TWNNM
  FROM XMLTable('$d/msg/rowOp/updateRow' PASSING CDCDOC as "d"
  COLUMNS
    "CUSTID" CHAR(5) PATH 'col[@name="CUSTID"]/char/afterVal',
    "STRTNM" VARCHAR(70) PATH 'col[@name="STRTNM"]/varchar/afterVal',
    "BLDGNB" VARCHAR(16) PATH 'col[@name="BLDGNB"]/varchar/afterVal',
    "PSTCD" VARCHAR(16) PATH 'col[@name="PSTCD"]/varchar/afterVal',
    "TWNNM" VARCHAR(35) PATH 'col[@name="TWNNM"]/varchar/afterVal'
  ) AS X ;

```

Note: Apply the PTF for APAR PM28385 (currently open) for the latest maintenance related to XMLTABLE function.

6.8.3 XML history objects

This section explores the use of pureXML to store historical data, and support temporal queries against that data. We consider an option to create and maintain pureXML data objects (representing a customer, for example) that contain a growing historical record of all the changes that have been applied to that object in the source systems.

Bitemporal data

Before proceeding to discuss XML support for temporal data, let us take note of another DB2 10 capability that supports the storage of historical data and the execution of temporal queries against that data: bitemporal data.

The concept of bitemporal data is that 4 additional columns (timestamp data type) are added to tabular data, reflecting the start and end points of 'system time' and 'business time'.

- ▶ System time is an auditable history of what the data looked like in the system at any point in history.
- ▶ Business time is an auditable history of the data, which reflects business corrections.

Bitemporal data support is a productivity feature of DB2 which can be used to maintain an auditable history of changes over time. Whenever a change is made to a table with bitemporal data support, the system time and the business time are both automatically tracked to maintain that auditable history.

Bitemporal data is not explained in any more detail in this book, because it is covered in other DB2 10 publications. For new DB2 systems, particularly for the relational data model, bitemporal data will probably be the most productive way of supporting temporal data, because the data structures and programming interface has all been built into DB2. Also, the use of bitemporal data in conjunction with XML is fully supported.

Temporal data with XML

XML also provides excellent support for temporal data.

The value of XML for historical data is broadly acknowledged as powerful and efficient. The paper "XML-Based Support for Database Histories and Document Versions" by Fusheng Wang, in 2004 describes an XML model for storing historical data, and supporting temporal queries. It is available at:

<http://wis.cs.ucla.edu/~wangfsh/publications/thesis.pdf>

Building XML History Objects from change data capture data streams

Most legacy data sources have been defined without system-provided temporal data support because bitemporal data is new in DB210 for z/OS.

Many transactional database systems are focussed almost exclusively on the current state of the data that they store, and may not have much support for historical data. This is a common scenario, which often leads to a decision to deploy a data warehouse (which is derived from the data in the transactional system) for the purpose of storing historical data and performing trend analysis to see how data has changed over time.

Change data capture and replication products are often used as a vehicle to feed changes to data warehouses. The change data capture data streams contain details of what the data changes were, and precisely when they happened, usually as a commit timestamp from the log record on the source system. This temporal reference data can be used by the ETL processes that maintain the data warehouse, in order to build a record of historical changes in the target data warehouse.

DB2 pureXML should be considered as a data type in the data warehouse for recording a history of changes to source data objects. pureXML would not be a good choice for data warehouse that seek to deliver OLAP-style structures. However, pureXML would be a good choice for other scenarios where the state of data at different points in the past is of interest. The following characteristics would tend to make pureXML attractive for historical data

- ▶ Large and sparse data structures could require extensive and costly relational database structures, but can be implemented with simplicity within an XML schema
- ▶ sub-document update allows changes replicated from source systems to be merged into a single XML document, combining all historical changes into a single XML document (as shown in Example 6-34 on page 122)
- ▶ SQL/XML provides powerful temporal query capabilities (provided the XML schema is designed to be friendly for temporal queries)
- ▶ XML and relational storage can co-exist in DB2's hybrid database environment, supporting queries with a mixture of SQL and XML expressions, and a mixture of relational and XML storage.

DB2 pureXML could act as a repository for historical data, with temporal query support as in Figure 6-11.

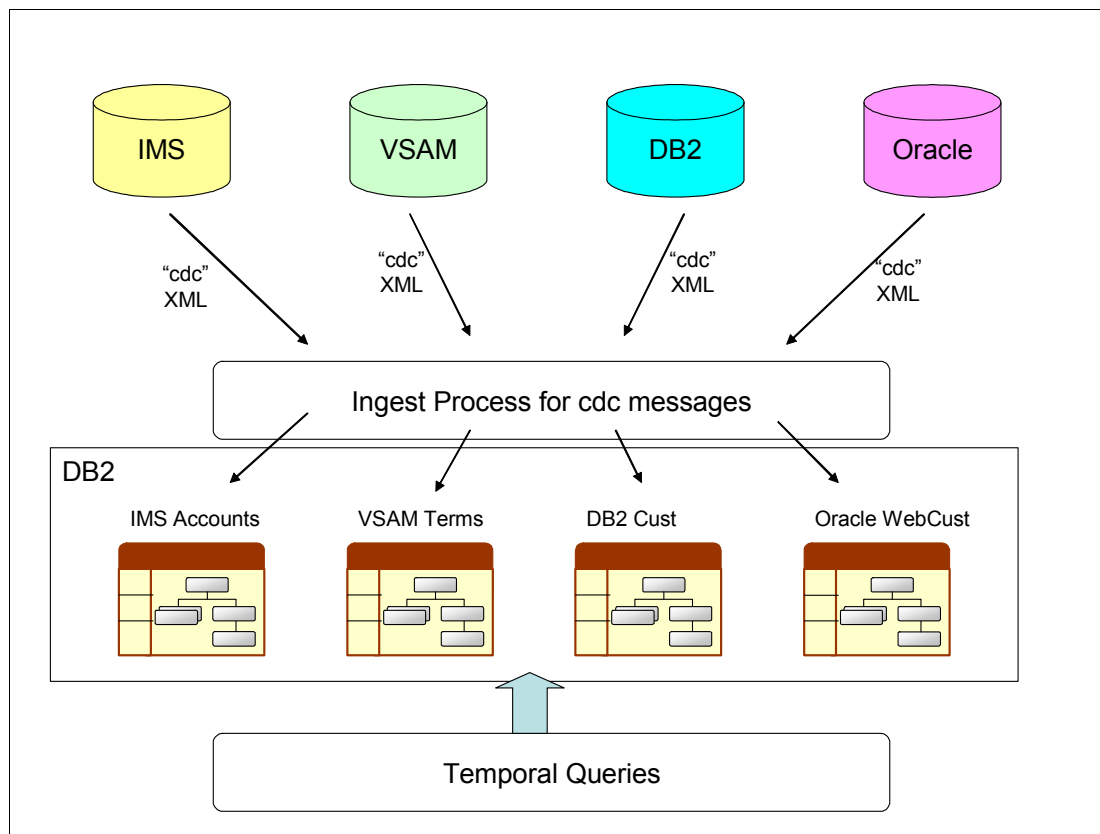


Figure 6-11 DB2 PureXML as historical repository



Using XML with Java

In Java applications, you can store XML data in DB2 databases or retrieve XML data from DB2 databases by using JDBC or SQLJ interface. Java also provides powerful APIs for XML processing, such as DOM, SAX and StAX.

In this chapter we show how JDBC applications handle XML data in DB2 for z/OS based on a scenario using the ISO20022 BankToCustomerStatement message.

This chapter contains the following:

- ▶ XML in Java
- ▶ The BankStmt application in Java
 - Setting up the environment
 - Insertion of rows with XML column values
 - Updates of XML columns
 - Retrieving XML data
 - Call stored procedure in Java
 - XSLT to transform XML document
 - Java interface to MQ

The description of the application scenario can be found in Chapter 3, “Application scenario” on page 45.

All Java examples are available for download as additional material as described in Appendix B, “Additional material” on page 273.

7.1 XML in Java

In DB2 tables, the XML built-in data type is used to store XML data in a column as a structured set of nodes in a tree format. You can write applications to store XML data in DB2 tables and retrieve XML data from tables.

The Java programming language and its database interface JDBC are very popular choices for XML application development. DB2 for z/OS provides a universal driver that supports both the JDBC and the SQLJ interface of the Java language. This driver is the IBM Data Server Driver for JDBC and SQLJ, also known as JCC (Java Common Client). It supports JDBC type 2 and type 4 driver and can connect to the DB2 family of products and Informix® Dynamic Server (IDS) database systems.

Two versions of the IBM Data Server Driver for JDBC and SQLJ are available.

- ▶ IBM Data Server Driver for JDBC and SQLJ version 3.x, JDBC 3.0-compliant.
- ▶ IBM Data Server Driver for JDBC and SQLJ version 4.x, JDBC 4.0-compliant.

We summarize the Java interfaces for XML data type support in Table 7-1.

Table 7-1 XML data type support in JCC3 and JCC4

JCC driver	JDBC support	Java interface for XML Data	Minimum Java level required	Jar file
JCC 3.x	JDBC 3.0	com.ibm.db2.jcc.DB2Xml	1.4	db2jcc.jar sqlj.jar
JCC 4.x	JDBC 4.0	java.sql.SQLXML	6.0	db2jcc4.jar sqlj4.jar

You control the level of JDBC support by specifying the appropriate set of files in the CLASSPATH. For example, if you want to use the JCC 4.x driver, include lib_dir/db2jcc4.jar and lib_dir/sqlj4.jar in your CLASSPATH.

In JDBC applications, you can:

- ▶ Register and remove XML schemas using Java methods.
- ▶ Create XML documents from application data using XML APIs.
- ▶ Store an entire XML document in an XML column using setXXX methods.
- ▶ Retrieve an entire XML document from an XML column using getXXX methods.
- ▶ Retrieve a sequence from a document in an XML column by using the SQL XMLQUERY function to retrieve the sequence into a serialized sequence in the database, and then using getXXX methods to retrieve the data into an application variable.
- ▶ Retrieve a sequence from a document in an XML column as a user-defined table by using the SQL XMLTABLE function to define the result table and retrieve it. Then use getXXX methods to retrieve the data from the result table into application variables.
- ▶ Invoke routines with XML parameters in Java applications.

JDBC3.0 provides some basic methods to retrieve XML document as Bytes, String or Stream, or set the XML document from Bytes, String or Stream. IBM Data Server Driver for JDBC and SQLJ 3.x also supports DB2Xml interfaces for XML processing.

JDBC4.0 introduces the SQLXML Java data type, which lets you map an XML data type table column to a Java data type. Use IBM Data Server Driver for JDBC and SQLJ 4.x, the latest

level. This driver provides the SQLXML interface and also support other DB2 features such as binary XML format. You have more flexibility in coding your application and better performance.

7.1.1 XML support in JDBC 3.0

When developing applications with JDBC3.0, you can use the Java standard interface `ResultSet` in order to retrieve XML data from DB2 tables. This interface provides getter methods for retrieving XML column values from the current row. Table 7-2 gives a a summary of the methods in the `ResultSet` interface.

Table 7-2 *JDBC 3.0 Getter methods of ResultSet*

Getter methods of ResultSet	Description
<code>getBytes</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as UTF-8 encoded bytes
<code>getString</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as a string in the Java programming language
<code>getAsciiStream</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as a stream of ASCII characters.
<code>getBinaryStream</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as a UTF-8 encoded binary stream
<code>getCharacterStream</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as a <code>java.io.Reader</code> object
<code>getObject</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as a <code>com.ibm.db2.jcc.DB2Xml</code> object

Note: The methods in Table 7-2 do not add an encoding declaration to the retrieved XML data.

The method `getObject` retrieves XML data into an object of type `DB2Xml`, which provides more getter method. `com.ibm.db2.jcc.DB2Xml` supports the methods shown in Table 7-3.

Table 7-3 *DB2Xml Getter Methods*

DB2Xml Getter Methods	Description
<code>getDB2Bytes()</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as UTF-8 encoded bytes
<code>getDB2XmlBytes()</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as a byte array in the Java programming language. The method converts the bytes to the target encoding and adds XML declaration with encoding tag
<code>getDB2String()</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as a string in the Java programming language
<code>getDB2XmlString()</code>	Retrieves the value of the designated column in the current row of this <code>ResultSet</code> object as a string in the Java programming language. The method will add XML declaration with encoding tag "ISO-10646-UCS-2"

DB2Xml Getter Methods	Description
getDB2AsciiStream()	Retrieves the value of the designated column in the current row of this ResultSet object as a stream of ASCII characters.
getDB2XmlAsciiStream()	Retrieves the value of the designated column in the current row of this ResultSet object as a stream of ASCII characters. The method will add XML declaration with encoding tag
getDB2BinaryStream()	Retrieves the value of the designated column in the current row of this ResultSet object as a UTF-8 encoded binary stream
getDB2XmlBinaryStream()	Retrieves the value of the designated column in the current row of this ResultSet object as a binary stream. The method converts the bytes to the target encoding and adds XML declaration with encoding tag
getDB2CharacterStream()	Retrieves the value of the designated column in the current row of this ResultSet object as a java.io.Reader object
getDB2XmlCharacterStream()	Retrieves the value of the designated column in the current row of this ResultSet object as a java.io.Reader object. The method will add XML declaration with encoding tag "ISO-10646-UCS-2"

The getDB2XmlXXX methods generate XML declarations with an encoding attribute for the retrieved XML data. For example, the methods getDB2String() and getDB2XmlString() return the XML data in the same encoding, USC-2, but the latter adds the appropriate encoding declaration to the XML document.

In a JDBC application, you can update or insert data into XML columns of a table at a DB2 data server using one of the setter methods of the interface PreparedStatement.

Following are the setXXX methods that are supported against XML columns in JDBC 3.0.

- ▶ setAsciiStream()
- ▶ setBinaryStream()
- ▶ setBlob()
- ▶ setBytes()
- ▶ setCharacterStream()
- ▶ setClob()
- ▶ setString()
- ▶ setObject(): supports DB2Xml, String, byte[], InputStream, Reader, CLOB, and BLOB as parameters.

7.1.2 XML support in JDBC 4.0

In JDBC 4.0, the main feature for XML support is the new SQLXML object, which is the mapping in the Java programming language for the SQL XML data type. The SQLXML interface provides methods for accessing the XML value as a String, a Reader or Writer, or as a Stream. The XML value may also be accessed through a Source or set as a Result, which are used with XML Parser APIs such as DOM, SAX, and StAX.

The interfaces ResultSet, CallableStatement, and PreparedStatement are enhanced with new getter or setter methods, such as ResultSet.getSQLXML to allow a programmer to

retrieve the value of the designated column of this `ResultSet` as a `java.sql.SQLXML` object in the Java programming language.

The `SQLXML` interface provides the methods listed in Table 7-4 for retrieving XML data from an `SQLXML` object.

Table 7-4 Methods to retrieve XML data from `SQLXML` object

SQLXML getter methods	Description
<code>getBinaryStream()</code>	Retrieves the XML value designated by this <code>SQLXML</code> instance as a stream
<code>getCharacterStream()</code>	Retrieves the XML value designated by this <code>SQLXML</code> instance as a <code>java.io.Reader</code> object
<code>getString()</code>	Returns a string representation of the XML value designated by this <code>SQLXML</code> instance
<code>getSource(Class<T> sourceClass)</code>	Returns a <code>Source</code> for reading the XML value designated by this <code>SQLXML</code> instance

`SQLXML` interface provides the methods listed in Table 7-5 for setting XML data to a `JDBC` object.

Table 7-5 Method to set XML value to `SQLXML` object

SQLXML setter methods	Description
<code>setBinaryStream()</code>	Retrieves a stream that can be used to write the XML value that this <code>SQLXML</code> instance represents
<code>setCharacterStream()</code>	Retrieves a stream to be used to write the XML value that this <code>SQLXML</code> instance represents
<code>setString()</code>	Sets the XML value designated by this <code>SQLXML</code> instance to the given <code>String</code> representation
<code>setResult(Class<T> resultClass)</code>	Returns a <code>Result</code> for setting the XML value designated by this <code>SQLXML</code> instance. void <code>setString(String value)</code>

For the implementation of the `BankStmt` application in Java, we mostly use `JDBC 4.0` standard `SQLXML` object in our programming.

7.1.3 Constructing XML document in Java

You are able to construct XML document using `DB2 SQL/XML` publishing functions. If your application holds information in application variables and wants to construct this data into an XML document, you can also do this by easily writing Java code with different XML APIs. In our example, we will construct a `GroupHeader` (subdocument of the `BankToCustomerStatement` message). We will use `DOM` APIs in this example because it is easy to understand. The `GroupHeader` we want to create is shown in Example 7-1.

Example 7-1 Creating a `GroupHeader` of the `BankToCustomerStatement` message

```
<GrpHdr>
  <MsgId>AAAASESS-FP-ACCR001</MsgId>
  <CreDtTm>2010-10-18T12:30:00+01:00</CreDtTm>
  <MsgPgntn>
```

```

    <PgNb>1</PgNb>
    <LastPgInd>true</LastPgInd>
  </MsgPgntn>
</GrpHdr>

```

The code for constructing XML as a DOM tree is shown in Example 7-2.

Example 7-2 Constructing XML as a DOM tree

```

import org.w3c.dom.Document;
import org.w3c.dom.Element;
import org.w3c.dom.Text;
...
String MsgIdStr="AAAASESS-FP-ACCR001";
String CreDtTmStr="2010-10-18T12:30:00+01:00";
String PgNbStr="1";
String LastPgIndStr="true";

DocumentBuilderFactory factory =
    DocumentBuilderFactory.newInstance();

try{
    DocumentBuilder builder = factory.newDocumentBuilder();
    Document document = builder.newDocument(); ❶

    //*****
    // Create element nodes
    //*****
    Element GrpHdr = document.createElement("GrpHdr"); ❷
    Element MsgId = document.createElement("MsgId");
    Element CreDtTm = document.createElement("CreDtTm");
    Element MsgPgntn = document.createElement("MsgPgntn");
    Element PgNb = document.createElement("PgNb");
    Element LastPgInd = document.createElement("LastPgInd");

    //*****
    // Create text nodes with designated values
    //*****
    Text MsgIdValue = document.createTextNode(MsgIdStr); ❸
    Text CreDtTmValue = document.createTextNode(CreDtTmStr);
    Text PgNbValue = document.createTextNode(PgNbStr);
    Text LastPgIndValue = document.createTextNode(LastPgIndStr);

    //*****
    // Append text node under element nodes
    //*****
    MsgId.appendChild(MsgIdValue); ❹
    CreDtTm.appendChild(CreDtTmValue);
    PgNb.appendChild(PgNbValue);
    LastPgInd.appendChild(LastPgIndValue);

    //*****
    // Construct the DOM tree
    //*****
    MsgPgntn.appendChild(PgNb); ❺
    MsgPgntn.appendChild(LastPgInd);

```

```
GrpHdr.appendChild(MsgId);
GrpHdr.appendChild(CreDtTm);
GrpHdr.appendChild(MsgPgntn);

document.appendChild(GrpHdr);

}catch (Exception e) {
    System.out.println("Some exception occur: " +
        e.getMessage());
}
```

Note:

1. Obtain a new instance of a DOM Document object to build a DOM tree
 2. Document.createElement() method to create element nodes with specified tagName
 3. Document.createTextNode() method to create text nodes with designated values
 4. Element.appendChild() method to append text node under element nodes
 5. Append element nodes and construct the DOM tree
-

7.1.4 Binary XML format in Java applications

DB2 10 for z/OS introduces a binary format for XML data to be used with Java applications, which is called Extensible Dynamic Binary XML DB2 Client/Server Binary XML Format (XDBX)¹. This format is an external representation of an XML value that is only used for exchange with a DB2 client application or the UNLOAD or LOAD utilities. The binary representation is smaller in size, and it saves the parsing cost.

The IBM Data Server Driver for JDBC and SQLJ can send XML data to the data server or retrieve XML data from the data server as textual XML data or binary XML data, as shown in Figure 7-1.

¹ See <http://www.ibm.com/support/docview.wss?uid=swg27019354&aid=1>.

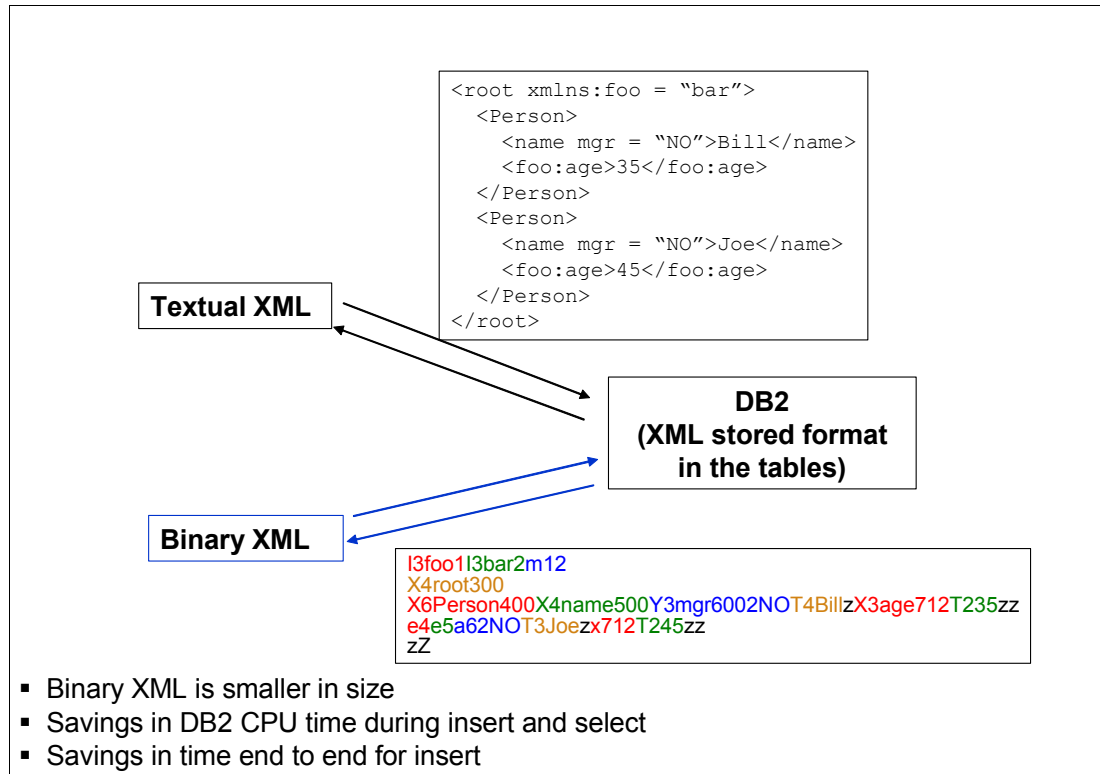


Figure 7-1 Exchange data as textual or binary XML format

You are able to set the JCC datasource property `xmlFormat` to control whether the data format is textual XML format or binary XML format. Possible values are:

- ▶ `com.ibm.db2.jcc.DB2BaseDataSource.XML_FORMAT_NOT_SET`
Specifies that binary XML format is used if the data server supports it. If the data server does not support binary XML format, textual XML format is used. This is the default.
- ▶ `com.ibm.db2.jcc.DB2BaseDataSource.XML_FORMAT_TEXTUAL`
Specifies that the XML textual format is used.
- ▶ `com.ibm.db2.jcc.DB2BaseDataSource.XML_FORMAT_BINARY`
Specifies that the binary XML format is used.

To change the XML format, you need set the datasource property first, then get a new connection, which will pick up your setting. Example 7-3 shows the Java code for setting the `xmlFormat`. Storage and retrieval of binary XML data requires the IBM Data Server Driver for JDBC and SQLJ version 4.9 or later. If you use the DB2 Client, V9.7 Fix Pack 3a and above provides this support.

Example 7-3 Setting the xmlFormat

```
DB2BaseDataSource db2ds=null;
...
db2ds.setXmlFormat(com.ibm.db2.jcc.DB2BaseDataSource.XML_FORMAT_BINARY);
con = ((DB2SimpleDataSource)db2ds).getConnection();
```

Note that the format of XML data is transparent to the application. The IBM Data Server Driver for JDBC and SQLJ presents binary XML data to the application only through the XML object interfaces. The user does not see the data in the binary XML format.

Accessing the XML value through SQLXML getSource() and setResult() methods in which the input or output data is in a non-textual representation can lead to improved processing performance. The SAX representation is the most efficient way to retrieve data that is in the binary XML format because the data does not undergo extra conversions from binary format to textual format.

7.2 The BankStmt application in Java

The BankToCustomerStatement, one of the ISO 20022 (Universal financial industry message scheme) message, is used to inform the account owner, or authorised party, of the entries booked to the account, and to provide the owner with balance information on the account at a given point in time.

The BankToCustomerStatement message is composed of two building blocks:

- ▶ **Group Header:** This building block is mandatory and present once. It contains elements such as Message Identification and CreationDateTime.
- ▶ **Statement:** This building block is mandatory and repetitive. It should be repeated for each account on which a statement is provided. The report contains components such as Balance and Entry.

In Chapter 3, “Application scenario” on page 45, we present our scenario to log and store the message for auditing purposes. The diagram in Figure 3-2 on page 49 illustrates the flow between the various code samples presented in this book.

In our Java application, we retrieve, manipulate, and re-save those XML documents in DB2, shred the previously saved XML message (see 7.2.5, “Call stored procedure to shred XML” on page 145) and query the message to produce an XML document to send out via an MQ Message. The display of this data is generated by combining the XML output with an XSLT file to produce an new HTML/XML display.

The purpose of the Java application is to demonstrate how to use Java with DB2 pureXML and we have chosen to just emphasize the different steps and choices made that are related to XML, disregarding irrelevant code and components.

7.2.1 Setting up the environment

Before creating the actual Java application, we assume that the necessary environment has been set up, which including

- ▶ Java SDK and IBM Data Server Driver for JDBC and SQLJ are installed on client side.
- ▶ The table is defined and the schema is registered on server side

Java environment on client

To run the BankStmt application, you need install an SDK for Java Version 6 or later. To verify the Java version you have, run the following command in your command windows.

```
java -version
```

To run the BankStmt application, you need setup the IBM Data Server Driver for JDBC and SQLJ version 4.9 or later on your client. If you use the DB2 Client, V9.7 Fix Pack 3a and above provides this driver. To verify the driver version you have, run the following command in your command windows.

```
java com.ibm.db2.jcc.DB2Jcc -version
```

Schema registration

We need to register XML schemas in the DB2 XML schema repository to allow us to validate our XML documents. To register the schema, you can either call the DB2-supplied stored procedures, or Command Line Processor, or via the Java method. Example 7-4 shows how to register the schema in a Java application by calling the registerDB2XmlSchema method.

Example 7-4 Register XML schema in Java application

```
String NameSchema = "CAMT053_JAVA";
String XSDFilename = "camt.053.001.02.xsd";
...
//*****
// Deregister the XML schema if it already exists.
//*****

try {
    ds.deregisterDB2XmlObject(
        "SYSXSR", NameSchema);
}
catch (SQLException e) {}

System.out.println("deregister complete");

//*****
//register the XML schema
//*****

fi[0]= new FileInputStream(XSDFilename);

// Schema Name Qualifiers, always this one, or blank which will default to SYSXSR
xmlSchemaNameQualifiers[0] = "SYSXSR";

// Schema Name
xmlSchemaNames[0] = NameSchema;

//Schema Location: Null means a schema can't be referred by "schemaLocation" in
document.
xmlSchemaLocations[0] = LocationName;

//The actual contents of the schema documents.
xmlSchemaDocuments[0] = new BufferedInputStream(fi[0]);

//Lengths of the actual contents of the schema documents.
xmlSchemaDocumentsLengths[0] = (int) fi[0].getChannel().size();

//Properties of the schema documents. Null means no properties.
xmlSchemaDocumentsProperties[0] = null;

//Lengths of the properties of the whole schema.
xmlSchemaDocumentsPropertiesLengths[0] = 0;

//Properties of the whole schema. Null means no properties.
xmlSchemaProperties = null;

//Lengths of the properties of the whole schema.
xmlSchemaPropertiesLength = 0;
```

```
System.out.println("Register one schema begins...");
```

```
ds.registerDB2XmlSchema(
    xmlSchemaNameQualifiers,
    xmlSchemaNames,
    xmlSchemaLocations,
    xmlSchemaDocuments,
    xmlSchemaDocumentsLengths,
    xmlSchemaDocumentsProperties,
    xmlSchemaDocumentsPropertiesLengths,
    xmlSchemaProperties,
    xmlSchemaPropertiesLength,
    false);
```

```
System.out.println("Register schema completed.");
```

DDL for tables in BankStmt application

The BankToCustomerStatement message is received from DB2 MQ Listener and stored in the BK_TO_CSTMR_STMT table.

The DDL for the table is shown in Example 7-5. The values of MSG_ID and MSG_CRE_DT_TM columns come from the XML document itself. The corresponding XML element name is MsgId (Message Identification), CreDtTm (Creation Date Time) in Group Header of XML document.

Example 7-5 DDL for table BK_TO_CSTMR_STMT

```
CREATE TABLE "BK_TO_CSTMR_STMT" (
    "MSG_ID" VARCHAR(35) WITH DEFAULT NULL ,
    "MSG_CRE_DT_TM" TIMESTAMP WITH TIMEZONE WITH DEFAULT NULL,
    "BK_TO_CSTMR_STMT" XML NOT NULL )
IN DATABASE XMLR5DB;
```

The XML documents can either be validated explicitly using the DSN_XMLVALIDATE built-in function, or you can automate XML schema validation by adding an XML type modifier to an XML column definition. This is discussed in Chapter 5, “Validating XML data” on page 73. In this chapter, we show the validation via DSN_XMLVALIDATE function.

7.2.2 Insertion of rows with XML column values

Our scenario, described in Chapter 3, “Application scenario” on page 45, shows the stored procedure inserting the XML document into the BK_TO_CUST_STMT table after validating and shredding out MSG_ID and MSG_CRE_DT_TM from the input parameter.

As Java provides various XML Parser APIs to parse XML documents, it is also easy to shred out some useful information from XML data and insert them into DB2.

Example 7-6 shows that we read the XML document from a file, then parse it and get the MSG_ID and MSG_CRE_DT_TM by using DOM APIs, then insert the data, and use DSN_XMLVALIDATE to validate the XML document during insert.

Example 7-6 Parsing XML value and inserting into DB2

```
SQLXML sq$xml1 = con1.createSQLXML();
```

```

OutputStream os = sqlxml1.setBinaryStream();
//*****
// get SQLXML object from file
//*****
try{
    fis = new FileInputStream(XMLFilename);
    int read;
    while ((read = fis.read ()) != -1) {
        os.write (read);
    }
} catch (FileNotFoundException e) {
    e.printStackTrace();
} catch (IOException e) {
    e.printStackTrace();
}

//*****
// parse the XML value with a DOM parser
//*****
try{
    DocumentBuilder parser =
        DocumentBuilderFactory.newInstance().newDocumentBuilder();
    Document result = parser.parse(XMLFilename); ❶

    //get MSG_ID
    NodeList nodes = result.getElementsByTagName("MsgId"); ❷
    MsgId = nodes.item(0).getFirstChild().getNodeValue(); ❸

    //get MSG_CRE_DT_TM
    nodes = result.getElementsByTagName("CreDtTm");
    CreDtTm = nodes.item(0).getFirstChild().getNodeValue();
    //convert XML timestamp format to java format
    CreDtTm = CreDtTm.replace("T", " "); ❹
} catch (Exception e) {
    e.printStackTrace();
}

//*****
// insert XML data, also use DSN_XMLVALIDATE to validate
//*****
String sql = "INSERT INTO BK_TO_CSTMTR_STMT " +
    "VALUES(?,?,DSN_XMLVALIDATE(CAST(? AS XML),'SYSXSR.CAMT053_JAVA'))"; ❺
pst = con1.prepareStatement(sql);

pst.setString(1, MsgId);
pst.setString(2, CreDtTm);
pst.setSQLXML(3, sqlxml1); ❻

pst.executeUpdate(); ❼

```

The numbered steps in Example 7-6 are explained as follows:

1. The `parser.parse()` method parses the content of the given file as an XML document and returns a new DOM Document object.

2. This step finds the element named “Msgld”, using the `getElementsByTagName()` method, and returns a nodelist.
3. Since we are only expecting 1 node in the nodelist, we just simply use `item(0)` to get the first node in the nodelist, and then use `getFirstChild().getNodeValue()` method to get the value of text node.
4. The XML datetime format use 'T' is a separator indicating that time-of-day, such as 2010-10-18T17:00:00+01:00, which is not compatible with Java datetime format, which require a space between day and hour, such as 2010-10-18 17:00:00+01:00. We just convert it here.
5. We call function `DSN_XMLVALIDATE(CAST(? AS XML))` to validate the XML document.
6. We set the parameter to the given `java.sql.SQLXML` object.
7. We execute the insert. The XML document is inserted as binary XML format (data that is in the Extensible Dynamic Binary XML DB2 Client/Server Binary XML Format), if the data server supports binary XML data.

7.2.3 Updates of XML columns

You can use the SQL UPDATE statement to update entire documents in an XML column, or update portions of XML documents using the XMLMODIFY function with a basic XPath updating expression.

To update the entire XML documents, you can execute a Java Statement or execute a `prepareStatement` with a `setXXX` method to set the designated parameter to an XML value.

To update portions of XML documents, use the SQL UPDATE statement with the XMLMODIFY built-in scalar function. The XMLMODIFY function specifies a basic updating expression that you can use to insert nodes, delete nodes, replace nodes, or replace the values of a node in XML documents that are stored in XML columns.

Assume that we have insert a `BankToCustomerStatement` message, and the statement that reports on booked entries and balances for a cash account is shown as Figure 7-2.

```

- <Stmt>
  <Id>AAAASESS-FP-STAT002</Id>
  <CreDtTm>2010-10-18T17:00:00+01:00</CreDtTm>
  + <FrToDt>
  + <Acct>
  - <Bal>
    - <Tp>
      - <CdOrPrtry>
        <Cd>OPBD</Cd>
        </CdOrPrtry>
      </Tp>
      <Amt Ccy="SEK">600000</Amt>
      <CdtDbtInd>CRDT</CdtDbtInd>
    + <Dt>
    </Bal>
  - <Bal>
    - <Tp>
      - <CdOrPrtry>
        <Cd>CLBD</Cd>
        </CdOrPrtry>
      </Tp>
      <Amt Ccy="SEK">399900</Amt>
      <CdtDbtInd>CRDT</CdtDbtInd>
    + <Dt>
    </Bal>
  - <Ntry>
    <Amt Ccy="SEK">200100</Amt>
    <CdtDbtInd>DBIT</CdtDbtInd>
    <Sts>BOOK</Sts>
    + <BookgDt>
    + <ValDt>
      <AcctSvcrRef>AAAASESS-FP-ACCR-02</AcctSvcrRef>
    + <BkTxCd>
    + <NtryDtls>
    </Ntry>
  </Stmt>

```

Figure 7-2 Bank To Customer Statement example

The first balance (Bal element with the Cd code OPBD) show that the book balance of the account at the beginning of the account reporting period is 600000 Swedish krona. There is one entry (Ntry) in the statement, and the code of CdtDbtInd is DBIT which indicates the balance is a debit balance, so the operation is a decrease. The Amount is 200100 Swedish krona. The second balance (Bal element with the Cd code CLBD) show the balance of the account at the end of the pre-agreed account reporting period. It is the sum of the opening booked balance at the beginning of the period and all entries booked to the account during the pre-agreed account reporting period, so the amount is 399900 Swedish krona.

Now imagine that we want to modify/correct the message, add a new entry which is a credit balance increase of 100100 Swedish krona. We need do the following modification.

1. Append a new entry (Ntry) after the first entry, record the credit balance of 100100 Swedish krona.
2. Modify the amount of ClosingBooked (CLBD) balance to 500000 Swedish krona

Example 7-7 shows how to modify an XML document, insert nodes and replace the values of a node.

Example 7-7 Modifying an XML document

```
//*****
```

```
// update XML document, add a new entry as the last entry
//*****

String sql = " UPDATE BK_TO_CSTMR_STMT "+
"SET BK_TO_CSTMR_STMT = XMLMODIFY ( "+ 1
" 'declare default element namespace " +
" '\urn:iso:std:iso:20022:tech:xsd:camt.053.001.02\'; "+
" insert nodes $newentry/newNtry/Ntry "+ 2
" after /Document/BkToCstmrStmt/Stmt/Ntry[fn:last()]', "+ 3
" XMLPARSE(DOCUMENT " +
"      '4
" /Document/BkToCstmrStmt/Bal[Tp/CdOrPrtry/Cd=\\"CLBD\"]/Amt " + 5
" with \\"500000\>' "+
"WHERE MSG_ID=? ";

pst = con1.prepareStatement(sql);
pst.setString(1, "AAAASESS-FP-STAT002");
pst.executeUpdate();
```

The numbered steps in Example 7-7 are as follows:

1. Invoke XMLMODIFY function to insert or replace XML nodes.
2. This UPDATE is to insert a new Ntry node.,
3. The fn:last function returns the number of Ntry nodes. /.../Ntry[fn:last()] is the last Ntry node of the document. The new Ntry node will be inserted after the last Ntry node

4. This UPDATE is to replace value of a node.
5. It replaces the Amt value when ../Cd is "CLBD".

7.2.4 Retrieving XML data

You can use one of the following ways to retrieve XML data:

- ▶ Use the `ResultSet.getSQLXML` method to retrieve the data. Then use a `SQLXML.getXXX` method to retrieve the data into a compatible output data type. This technique requires JDBC 4.0 or later.
- ▶ Use a `ResultSet.getXXX` method other than `ResultSet.getObject` to retrieve the data into a compatible data type.
- ▶ If you use JCC3 driver that don't support JDBC 4.0, you can use the `ResultSet.getObject` method to retrieve the data, and then cast it to the `DB2Xml` type and assign it to a `DB2Xml` object. Use a `DB2Xml.getDB2XXX` or `DB2Xml.getDB2XmlXXX` method to retrieve the data into a compatible output data type.

You can retrieve the entire XML document or partial XML document using `XMLQUERY` function. You can retrieve data from XML columns in a table as binary XML data (data that is in the Extensible Dynamic Binary XML DB2 Client/Server Binary XML Format), if the data server supports binary XML data. Example 7-8 shows how to retrieve the entire or partial XML document.

Example 7-8 Retrieving the entire or partial XML document

```
//*****
// retrieve XML document to a SQLXML object
//*****
String sql = "SELECT BK_TO_CSTMR_STMT FROM BK_TO_CSTMR_STMT WHERE MSG_ID = ?";
pst = con1.prepareStatement(sql);
pst.setString(1, "AAAASESS-FP-STAT002");
ResultSet rs=pst.executeQuery();

while (rs.next()) {
    sqlxml=rs.getSQLXML(1);
    actualResults.println(sqlxml.getString());
}

//*****
// retrieve the XML nodes by using XMLQUERY to a SQLXML object
//*****
sql = " SELECT XMLQUERY(" +
    "'declare default element namespace " +
    "\"urn:iso:std:iso:20022:tech:xsd:camt.053.001.02\""; " +
    "/Document/BkToCstmrStmt/Stmt/Bal/Amt' " +
    "passing BK_TO_CSTMR_STMT ) " +
    "FROM BK_TO_CSTMR_STMT " +
    "WHERE MSG_ID = ?";
pst = con1.prepareStatement(sql);
pst.setString(1, "AAAASESS-FP-STAT002");
rs=pst.executeQuery();

while (rs.next()) {
    sqlxml=rs.getSQLXML(1);
    actualResults.println(sqlxml.getString());
}
```

 }

7.2.5 Call stored procedure to shred XML

The BankStmt application will continue to shred the BankToCustomerStatement message from BK_TO_CSTMTR_STMT table into a STMT table, which demonstrates a simple hybrid design, and how to populate using the INSERT from SELECT with XMLTABLE. We plan to do the shredding in a native stored procedure. Our java application will retrieve the XML message and call the stored procedure with the XML value as parameter.

Example 7-9 shows the definition for STMT table, each row contains one statement (Stmt building block) of the BankToCustomerStatement message, e.g. STMT_ID column is corresponding to the "Id" node under Stmt element, and STMT_XML column will be the sub document (the Stmt building block) of the BankToCustomerStatement message.

Example 7-9 DDL for STMT table

```
CREATE TABLE "STMT" (
  "STMT_ID" VARCHAR(35) NOT NULL ,
  "MSG_ID" VARCHAR(35) ,
  "MSG_CRE_DT_TM" TIMESTAMP ,
  "ELECTRNC_SEQ_NB" BIGINT ,
  "LGL_SEQ_NB" BIGINT ,
  "CRE_DT_TM" TIMESTAMP NOT NULL ,
  "FR_DT_TM" TIMESTAMP ,
  "TO_DT_TM" TIMESTAMP ,
  "RPTG_SRC_CD" CHAR(4) NOT NULL ,
  "RPTG_SRC_PRTRY" VARCHAR(35) NOT NULL ,
  "ADDTL_INF" VARCHAR(140) NOT NULL,
  "STMT_XML" XML NOT NULL )
IN DATABASE XMLR5DB ;
```

Example 7-10 shows our code for creating the SQL procedure in Java.

Example 7-10 Creating a SQL stored procedure

```
stmt.executeUpdate(" +
  "CREATE PROCEDURE MYSP(IN parm1 XML,OUT parm2 XML)" + ❶
  "LANGUAGE SQL          " +
  "APPLICATION ENCODING SCHEME UNICODE" +
  "DISABLE DEBUG MODE " +
  "BEGIN                  " +
  "DECLARE var1 XML;     " +
  "SET var1 = parm1;     " +
  "INSERT INTO STMT(     " + ❷
  "  STMT_ID,            " +
  "  MSG_ID,             " +
  "  MSG_CRE_DT_TM,      " +
  "  ELECTRNC_SEQ_NB,   " +
  "  LGL_SEQ_NB,         " +
  "  CRE_DT_TM,          " +
  "  FR_DT_TM,           " +
  "  TO_DT_TM,           " +
  "  RPTG_SRC_CD,        " +
  "  RPTG_SRC_PRTRY,     " +
```

```

" ADDTL_INF,          " +
" STMT_XML           " +
")                  " +
"SELECT T.STMT_ID,   " +
"   T.MSG_ID,       " +
"   T.MSG_CRE_DT_TM," +
"   T.ELECTRNC_SEQ_NB," +
"   T.LGL_SEQ_NB,   " +
"   T.CRE_DT_TM,    " +
"   T.FR_DT_TM,     " +
"   T.TO_DT_TM,     " +
"   COALESCE(T.RPTG_SRC_CD,'') AS PRTG_SRC_CD," + 3
"   COALESCE(T.RPTG_SRC_PRTRY,'') AS RPTG_SRC_PRTRY," +
"   COALESCE(T.ADDTL_INF,'') AS ADDTL_INF," +
"   XMLDOCUMENT(T.STMT_XML)" + 4
"FROM XMLTABLE(     " +
"  XMLNAMESPACES(DEFAULT " +
"    'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02')," +
"  '$var1/Document/BkToCstmrStmnt/Stmnt'" +
"  PASSING var1 as \"var1\"" +
"  COLUMNS STMT_ID VARCHAR(35) PATH 'Id'," +
"            MSG_ID VARCHAR(35)  PATH '../GrpHdr/MsgId'," +
"            MSG_CRE_DT_TM TIMESTAMP PATH '../GrpHdr/CreDtTm'," +
"            ELECTRNC_SEQ_NB BIGINT PATH 'ElctrncSeqNb'," +
"            LGL_SEQ_NB BIGINT    PATH 'LglSeqNb'," +
"            CRE_DT_TM  TIMESTAMP PATH 'CreDtTm'," +
"            FR_DT_TM  TIMESTAMP PATH  'FrToDt/FrDtTm'," +
"            TO_DT_TM  TIMESTAMP PATH  'FrToDt/ToDtTm'," +
"            RPTG_SRC_CD  CHAR(4) PATH  'RptgSrc/Cd'," +
"            RPTG_SRC_PRTRY VARCHAR(35) PATH 'RptgSrc/Prtry'," +
"            ADDTL_INF  VARCHAR(144) PATH  'AddtlStmntInf'," +
"            STMT_XML  XML    PATH    '.'" +
"  )AS T;          " +
//*****
// For the output parameter, You can do some more XML operations,
// here we just simply set the output parameter the same as input
//*****
"SET parm2 = var1; " +
"END
);

```

Example 7-10 demonstrates that:

1. In DB2 10, we can use XML as the data type for a parameter of a native SQL procedure, and an XML SQL variable declared within the procedure.
2. We use INSERT from SELECT with XMLTABLE function for the shredding. If the BankToCustomerStatement message includes multiple statements, we are able to break one message into multiple statements (one rows for each statement) into the STMT table by using the XMLTABLE function.
3. To insert a nullable value into a NOT NULL column, we use the COALESCE function.
4. When inserting the XML value from XMLTABLE, we need to use the XMLDOCUMENT function to return a constructed XML value

To call the stored procedures in Java, you should invoke methods in the `CallableStatement` class. The basic steps for calling a stored procedure using standard `CallableStatement` methods are:

1. Invoke the `Connection.prepareCall` method with the `CALL` statement as its argument to create a `CallableStatement` object.
2. Invoke the `CallableStatement.setXXX` methods to pass values to the input parameters (parameters that are defined as `IN` or `INOUT` in the `CREATE PROCEDURE` statement).
3. Invoke the `CallableStatement.registerOutParameter` method to register parameters that are defined as `OUT` in the `CREATE PROCEDURE` statement with specific data types.
4. Invoke the `CallableStatement.executeXXX` methods to call the stored procedure.
5. Invoke the `CallableStatement.getXXX` methods to retrieve values from the `OUT` parameters or `INOUT` parameters.

The sample code to prepare, call the stored procedure, and retrieve the XML data from `OUT` parameters is shown in Example 7-11.

Example 7-11 Handling the SQL stored procedure

```
String sql = "CALL MYSP(?,?)";
CallableStatement cstmt = con1.prepareCall(sql);

//initialize the parms
SQLXML xml1 = con1.createSQLXML();

stmt = con1.createStatement();
sql = "SELECT BK_TO_CSTMRT_STMT FROM BK_TO_CSTMRT_STMT " +
      "WHERE MSG_ID='AAAASESS-FP-STAT003-4'";
ResultSet rs = stmt.executeQuery(sql);
if(rs.next())
    xml1=rs.getSQLXML(1);

actualResults.println("value of input XML:");
actualResults.println(xml1.getString());

cstmt.setSQLXML(1, xml1);
cstmt.registerOutParameter(2, java.sql.Types.SQLXML);
cstmt.execute();

xml1 = cstmt.getSQLXML(2);
actualResults.println("value of output XML:");
actualResults.println(xml1.getString());

cstmt.close();
```

7.2.6 XSLT to transform XML document

With XSLT (see 1.2.4, “Extensible Stylesheet Language” on page 11) you can transform an XML document into another XML document, or another type of document that is recognized by a browser, like HTML and XHTML. With XSLT you can add or remove elements and attributes to or from the output file. You can also rearrange and sort elements, perform tests and make decisions about which elements to hide and display.

JAXP (Java API for XML Processing) provides interfaces in package javax.xml.transform allowing applications to invoke an XSLT transformation.

In the BankStmt application, we want to generate a report that lists all Entries which have an amount greater than 100,000 SEK. The expected output is shown in Example 7-12. In the output, we list the account Id, amount, date time, and account servicer reference of the entry.

Example 7-12 Expecting output after transform

```
<?xml version="1.0" encoding="UTF-8"?>
<Result xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02"
  xmlns:iso20022="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02">
  <Acct Id="50000000054910000005">
    <Ntry>
      <Amt Ccy="SEK">150000</Amt>
      <DtTm>2010-10-19T10:15:00+01:00</DtTm>
      <AcctSvcrRef>AAAASESS-FP-ACCR-03</AcctSvcrRef>
    </Ntry>
  </Acct>
  <Acct Id="50000000054910000006">
    <Ntry>
      <Amt Ccy="SEK">200000</Amt>
      <DtTm>2010-10-20T10:15:00+01:00</DtTm>
      <AcctSvcrRef>AAAASESS-FP-ACCR-05</AcctSvcrRef>
    </Ntry>
  </Acct>
</Result>
```

We use an XSLT file to transform the original BankToCustomerStatement message to a new XML document. The simple XSLT file is shown in Example 7-13.

Example 7-13 XSLT example to transform from XML to XML

```
<xsl:stylesheet version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:iso20022="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02"
  xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02">
  <!-- root template -->
  <xsl:template match="/">
  <Result>
    <xsl:for-each 1
      select="/iso20022:Document/iso20022:BkToCstmrStmt/iso20022:Stmt/iso20022:Ntry">
      <xsl:apply-templates select="."/> 2
    </xsl:for-each>
  </Result>
</xsl:template>
  <!-- Ntry template -->
  <xsl:template match="iso20022:Ntry">
    <xsl:if test="iso20022:Amt>100000"> 3
    <Acct Id="{../iso20022:Acct/iso20022:Id/iso20022:Othr/iso20022:Id}"> 4
    <Ntry>
      <Amt Ccy="{iso20022:Amt/@Ccy}">
        <xsl:value-of select="iso20022:Amt"/>
      </Amt>
      <DtTm>
        <xsl:value-of select="iso20022:BookgDt/iso20022:DtTm"/>
```

```

        </DtTm>
        <AcctSvcrRef>
            <xsl:value-of select="iso20022:AcctSvcrRef"/>
        </AcctSvcrRef>
    </Ntry>
</Acct>
</xsl:if>
</xsl:template>
</xsl:stylesheet>

```

The numbered steps in Example 7-13 are:

1. The xsl:for-each element repeat for each Ntry node, which selected by the XPath `/iso20022:Document/iso20022:BkToCstmrStmnt/iso20022:Stmnt/iso20022:Ntry`
2. For each Ntry node, apply Ntry template to do more transform
3. In Ntry template, the xsl:if element test if the Amt>100000
4. Construct the expecting Acct element with attribute Id

For more information about XSLT standard, reference <http://www.w3.org/TR/xslt>.

The Java application reads the BankToCustomerStatement message stored in BK_TO_CSTMTR_STMT table, then transforms it into a new XML document based on the XSLT file and stores it into a file. The code snippet is shown in Example 7-14.

Example 7-14 Java application to transform XML document

```

//retrieve xml document from BK_TO_CSTMTR_STMT
String sql = "SELECT BK_TO_CSTMTR_STMT FROM BK_TO_CSTMTR_STMT " +
            "WHERE MSG_ID='AAAASESS-FP-STAT003-4'";
ResultSet rs = stmt.executeQuery(sql);
if(rs.next())
    xml1=rs.getSQLXML(1);

// JAXP reads data using the Source interface
Source xmlSource = new StreamSource(xml1.getBinaryStream());
Source xsltSource = new StreamSource(xsltFile);

// the factory pattern supports different XSLT processors
TransformerFactory transFact =
TransformerFactory.newInstance();
Transformer trans = transFact.newTransformer(xsltSource);

//transform and put the new XML document into a file
Result result = new StreamResult(xmloutputFile);
trans.transform(xmlSource, result);

```

Transforming an XML document to a HTML display is quite similar. The only difference is that we need to transform the HTML document with standard tag and attribute name. Example 7-15 i shows that.

Example 7-15 XSLT example to transform from XML to HTML

```

<xsl:stylesheet version="1.0"
    xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
    xmlns:iso20022="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02"
>

```

```

<!-- root template -->
<xsl:template match="/">
<html> ❶
  <body>
    <table border="1" cellspacing="0"> ❷
      <tr>
        <th>Acct(Id)</th> ❸
        <th>Amt(Ccy="SEK")</th>
        <th>DtTm</th>
        <th>AcctSvcrRef</th>
      </tr>
      <xsl:for-each select="//iso20022:Ntry"> ❹
        <xsl:apply-templates select="."/>
      </xsl:for-each>
    </table>
  </body>
</html>
</xsl:template>
<!-- Ntry template -->
<xsl:template match="iso20022:Ntry">
  <xsl:if test="iso20022:Amt>100000">
    <tr>
      <td>
        <xsl:value-of
select="../iso20022:Acct/iso20022:Id/iso20022:0thr/iso20022:Id"/>
      </td>
      <td>
        <xsl:value-of select="iso20022:Amt"/>
      </td>
      <td>
        <xsl:value-of select="iso20022:BookgDt/iso20022:DtTm"/>
      </td>
      <td>
        <xsl:value-of select="iso20022:AcctSvcrRef"/>
      </td>
    </tr>
  </xsl:if>
</xsl:template>
</xsl:stylesheet>

```

Notes:

1. Start building a html document
2. Create a simple table in the HTML document
3. Define the header information using <th> tag
4. For each Ntry, apply Ntry template to convert it into one row in the table

The display of the HTML document is in Figure 7-3.

Acct (Id)	Amt (Ccy="SEK")	DtTm	AcctSvcrRef
50000000054910000005	150000	2010-10-19T10:15:00+01:00	AAAASESS-FP-ACCR-03
500000000054910000006	200000	2010-10-20T10:15:00+01:00	AAAASESS-FP-ACCR-05

Figure 7-3 HTML output after XSLT

7.2.7 Java interface to MQ

WebSphere MQ allows you to easily exchange information across different platforms, integrating new and existing business applications. The BankStmt application scenario (XML message logging and auditing, as described in Chapter 3, “Application scenario” on page 45) describes a common situation, where XML messages are flowing over a WebSphere MQ. In Chapter 6, “DB2 SQL/XML programming” on page 87 we described how to capture these messages and store them in DB2 pureXML for auditing purposes. In this section, after retrieving the messages from DB2 and doing some transformations using XSLT, our Java application puts the generated XML document in the WebSphere MQ.

To put a message into a queue, you need to connect to the queue manager first, which provides queuing services to the applications, then open a queue and put the message into it. Figure 7-4 show the process.

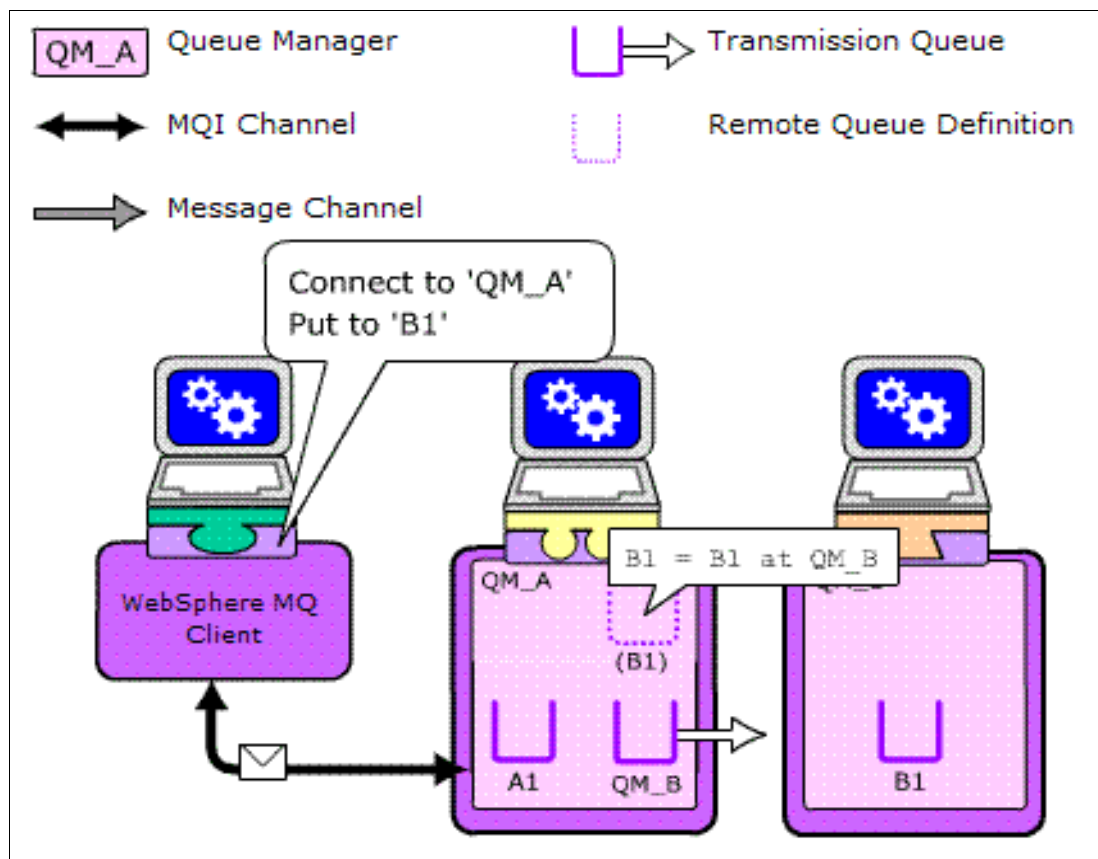


Figure 7-4 Put a message into a Queue

WebSphere MQ classes for Java Message Service (also referred to as WebSphere MQ JMS) is a set of Java classes that implement Oracle Sun’s Java Message Service (JMS) interfaces to enable JMS programs to access WebSphere MQ systems.

Using WebSphere MQ JMS as the API to write WebSphere MQ applications has a number of benefits. Some advantages derive from JMS being an open standard with multiple implementations. Other advantages come from additional features that are present in WebSphere MQ JMS, but not in WebSphere MQ base Java, such as asynchronous message delivery, message selectors, support for publish/subscribe messaging. WebSphere MQ JMS APIs can be used to connect to MQ on multiple platforms.

The example code to put the message into a queue is shown in Example 7-16.

Example 7-16 Java example to put a message into a queue

```
//*****
// Before starting the example, you need input the following information:
// IP address(ipaddress), port number(port), channel name(channel)
// queue manager name(queueMgr) and queue name(queueName)
//*****

//*****
// Maps the specified key to the specified value into properties
//*****
properties = new Hashtable();

properties.put("hostname", ipAddress);
properties.put("port", new Integer(Integer.parseInt(port)));
properties.put("channel", channel);

//*****
// Connect to MQ
//*****
try {
    queueManager = new MQQueueManager(queueMgr, properties);
}
catch (java.lang.NoClassDefFoundError e) {
    ...
}

//*****
// Open a queue
//*****
MQQueue system_default_local_queue=null;

try {
    system_default_local_queue = queueManager.accessQueue(queueName,
        MQConstants.MQOO_INPUT_AS_Q_DEF | MQConstants.MQOO_OUTPUT,
        null, null, null);
}
catch (MQException ex) {
    ...
}

//*****
// Put a message into the queue
//*****
MQMessage put_msg = new MQMessage();
put_msg.characterSet = 1208;
try {
    put_msg.writeUTF(xmlmessage);
}
catch (Exception ex) {
    ...
}

try {
```



```
    system_default_local_queue.put(put_msg, new MQPutMessageOptions());  
  }  
  catch (MQException ex) {  
    ...  
  }
```

All Java example are available as additional material as described in Appendix B, “Additional material” on page 273.

The sample code includes:

- ▶ RegisterSchema.java
This class registers the XML schema to the DB2 databases
- ▶ InsertXML.java
This class validate and insert XML data into db2 table.
- ▶ UpdateXML.java
This class demo partial updates of XML documents
- ▶ SelectXML.java
This class demo the retrieving entire or partial XML document to a SQLXML object
- ▶ XMLProcedure.java
This class creates SQL stored procedure with XML as parameter to shred XML document, then call the SQL stored procedure from Java
- ▶ TransformXML.java
This class transform XML document retrieved into an new XML or HTML document
- ▶ SendMQMessage.java
This class send a XML message generated by XSLT to WebSphere MQ

We list the data we used as below:

- ▶ Schema of bank to customer statement message
camt.053.001.02.xsd
- ▶ XML file of bank to customer statement message
camt.053.001.02.xml
camt.053.001.03.xml
camt.053.001.04.xml
- ▶ XSLT file to transform XML message
camt.053.001.04.xsl
camt.053.001.05.xsl
- ▶ XML document send to WebSphere MQ
xmloutput.xml



Using XML with COBOL

With Enterprise COBOL for z/OS® V4.1, you can integrate COBOL and Web-based business processes in Web services, XML, Java™, and COBOL applications. This interoperability enables you to capitalize on existing IT investment while incorporating new, Web-based applications as part of your organization's infrastructure.

In this chapter we demonstrate the COBOL support for the DB2 pureXML format by implementing a small COBOL application based on the BankToCustomerStatement message of the ISO20022 standard presented in Chapter 3, “Application scenario” on page 45.

We also briefly touch on the native COBOL facilities for XML and discuss how they complement those of DB2 pureXML.

The chapter contains the following:

- ▶ XML representation in COBOL
- ▶ The BankStmnt application in COBOL
- ▶ COBOL functions for manipulating XML

The complete source code, DDL, XML schema, and scripts used in the application are made available for download as described in Appendix B, “Additional material” on page 273.

8.1 XML representation in COBOL

COBOL offers different options for working with XML data. Most importantly, there is support for the pureXML data type in DB2 in a couple of different variations, but in some cases ordinary binary or character-based types may also be used. In addition, the file reference variable was applicable to LOBs and XML.

As XML is always stored in Unicode, special attention has to be paid to which code pages are used in the application and how to avoid data conversion.

We use the DB2 precompiler throughout this chapter. Experiences may vary slightly if using the co-processor.

8.1.1 XML host variables in COBOL

In DB2, the data type XML is a basic data type with its own representation and associated simple functions. You can insert and modify data as XML, and you can retrieve data as XML. In COBOL, the XML data type always builds on one of the existing LOB formats. XML variable declarations built on the basic LOB types are shown in Example 8-1.

Example 8-1 XML host variables in COBOL

```
01 DOC-AS-CLOB    IS SQL TYPE IS XML AS CLOB(100K).
01 DOC-AS-BLOB    IS SQL TYPE IS XML AS BLOB(100K).
01 DOC-AS-DBCLOB IS SQL TYPE IS XML AS DBCLOB(100K).
```

These declarations are translated by the precompiler as shown in Example 8-2.

Example 8-2 XML host variables after transformation by the DB2 pre-compiler

```
01 DOC-AS-CLOB.
   49 DOC-AS-CLOB-LENGTH          PIC S9(9) COMP-5.
   49 DOC-AS-CLOB-DATA            PIC X(102400).
01 DOC-AS-BLOB.
   49 DOC-AS-BLOB-LENGTH          PIC S9(9) COMP-5.
   49 DOC-AS-BLOB-DATA            PIC X(102400).
01 DOC-AS-DBCLOB.
   49 DOC-AS-DBCLOB-LENGTH        PIC S9(9) COMP-5.
   49 DOC-AS-DBCLOB-DATA          PIC G(102400) DISPLAY-1.
```

Whether you want to use CLOBs, DBCLOBs or BLOBs as the base format depends on how and whether you want to manipulate the contents of the XML file in the application.

One significant difference between the base format of BLOB compared to that of CLOB (or DBCLOB) is the encoding. XML is always stored in UTF-8 Unicode whereas the COBOL application will work in EBCDIC or UTF-16 Unicode, so data conversion and data encoding is always an issue you need to consider.

Data encoding

The character-based formats are what are referred to as *externally encoded*, whereas the binary-based formats are referred to as *internally encoded*. A variable with subtype FOR BIT DATA is also considered internally encoded.

Externally encoded means that the code page is derived from the CCSID of the host variable, if one is specified, or else from the application code page. Internally encoded means that the code page is derived from the data itself.

An XML document may or may not contain an encoding declaration, regardless of whether it is internally or externally encoded. It is placed as an attribute of the XML declaration as can be seen in Example 8-3.

Example 8-3 XML declaration with encoding declaration as an attribute

```
<?xml version="1.0" encoding="IBM037"?>
```

However, the significance of the encoding declaration is different depending on the format. For internally encoded documents, the encoding declaration is a factor in deciding the code page of the XML document, together with the Unicode Byte Order Mark (BOM) and the rest of the XML declaration. The BOM is an optional byte sequence preceding the XML document denoting the byte order of the following text, it is used for Unicode documents only.

The exact algorithm for determining the code page of an internally encoded XML document is as follows:

- ▶ If the document contains a BOM, the BOM decides the Endianness. Possible outcomes are UTF-16, UTF-32, Big Endian or Little Endian.
- ▶ If no BOM is present, and the XML document contains an encoding declaration, the encoding declaration decides the code page. Possible outcome is any valid code page, it does not have to be UTF-n encoding.
- ▶ If no BOM and no encoding declaration is present, and the document contains an XML declaration, the code page is UTF-8 if the document is written in single-byte ASCII, UTF-16 if the document is double-byte ASCII.
- ▶ If no BOM and no XML declaration is present, the code page is UTF-8.

Note: If there is a mismatch in Endianness between BOM and the UTF-16 XML document, the insert or update fails.

For externally encoded documents any encoding declarations are ignored by DB2, instead the code page is determined as with any character data. The default is the application code page, but you may override it in the SQLDA, or by an explicit variable declaration as shown in Example 8-4. This is probably the easiest way, but it requires that the SQLDA is not in use for the variables in question.

Example 8-4 Explicit declaration of variable CCSID

```
01 XMLVAR USAGE IS SQL TYPE IS XML AS CLOB
   EXEC SQL DECLARE :XMLVAR VARIABLE CCSID 277 END-EXEC.
```

We conclude that the internally encoded variables and externally encoded variables give different results when used for inserting XML documents into DB2. This is summarized in Table 8-1 on page 158 where the effect of inserting various XML documents using either XML AS BLOB or XML AS CLOB is shown.

It is assumed that the XML AS CLOB variable has an EBCDIC CCSID associated with it, either through explicit declaration or from the application code page. It is also assumed that the EBCDIC code page referred to is the same through the table. The XML AS BLOB, of course, has no associated CCSID.

Table 8-1 Insert an XML document with different host variable types

Document encoding	XML AS CLOB, CCSID is EBCDIC	XML AS BLOB
EBCDIC document with EBCDIC encoding declaration	Insert successful	Insert successful
EBCDIC document with UTF-8 encoding declaration	Insert successful	SQL error -20398
UTF-8 document with EBCDIC encoding declaration	SQL error -20398	SQL error -20398
UTF-8 document with UTF-8 declaration	SQL error -20398	Insert successful

Regardless of the encoding format used to enter XML data into DB2, care should be taken to ensure that the encoding declaration is accurate, as data may later be transmitted to other components or applications relying on the encoding declaration. If the encoding is not accurate, errors may occur.

For details on how DB2 deals with with mixed data in character string, see:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db2z10.doc.sqlref/db2z_mixeddatainchar.htm

Avoiding data conversion

In general it is recommended to avoid data conversion as this is both costly in terms of CPU usage, and also has the potential for data loss. Data loss can occur if the resulting code page does not hold all the code points present in the original document.

When an XML document is inserted into DB2, data conversion always takes place if the code page of the XML document is anything but UTF-8. This is the case whether the variable holding the XML document is internally or externally encoded, so for the simple case of inserting an XML document in EBCDIC from a COBOL application, data conversion cannot be avoided even if using a BLOB as the base for the XML column.

On the other hand, when retrieving an XML document from DB2 data conversion only occurs if using an externally encoded format for the target variable and if the associated code page is different from UTF-8. So here data conversion is avoided when using a binary format, but then you have the data in Unicode which may or may not be desirable.

For COBOL applications that work with XML documents in an EBCDIC code page, either because they are received and transmitted to other applications in that format, or because they are created and manipulated entirely in COBOL, a good choice would be to use externally encoded variables because that ensures that the necessary data conversions are performed.

However, if the COBOL application receives the XML data from, or transmits XML data to an application using another code page, e.g. via WebSphere MQ, there is the risk of performing an unnecessary interim data conversion.

In this case it might be an idea to use an internally encoded format if the external application either uses Unicode, or is able to perform the conversion from Unicode to its own code page. The other option is using external encoding with the CCSID of the external application. Both these choices eliminate the need for an interim conversion of the data, but both require that the COBOL application can handle the XML data in the other application's code page.

In Figure 8-1 it is shown how the interim code conversion can be avoided by using an internally encoded variable for the XML data.

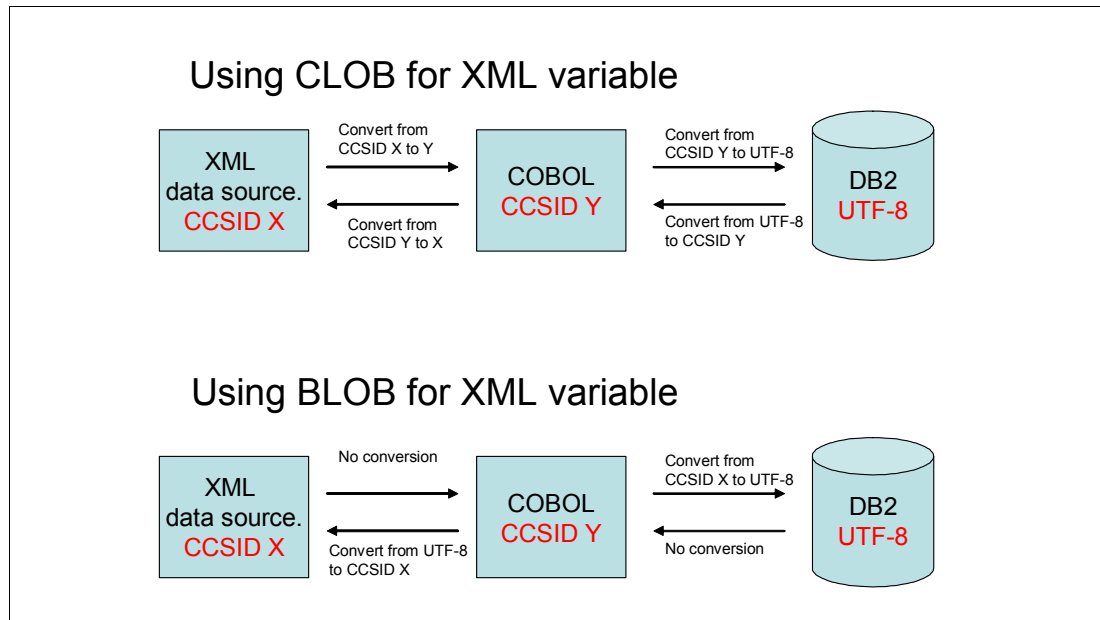


Figure 8-1 Data conversion in a three-layer structure using CLOBs or BLOBs

8.1.2 Using non-XML variables for XML data

In general, it is recommended to use XML variables for XML data but in some cases it might be practical to consider other options. One such case could be when developing COBOL stored procedures. At present, the XML data type is not supported for external stored procedures.

It is possible to use non-XML variables for insert and update, and for retrieval of XML documents. Possible data types are CLOB, BLOB, DBCLOB, CHAR, VARCHAR, GRAPHIC, VARGRAPHIC, BINARY AND VARBINARY.

You can either let DB2 handle the conversion implicitly, or explicitly call the conversion functions. XMLSERIALIZE converts the host variable type to XML whereas XMLPARSE converts XML to a host variable type. It is recommended to let DB2 perform the conversion implicitly whenever possible, as this is more efficient.

In more complex SQL/XML, calls, e.g. XMLMODIFY, it is not possible to substitute the XML type with other host variable types. In this case it might be necessary to perform an explicit conversion by calling XMLPARSE or XMLSERIALIZE.

Note: Always use XMLSERIALIZE WITHOUT XMLDECLARATION with an externally encoded host variable.

The XML declaration resulting from an explicit serialization always contains encoding="UTF-8" because the conversion is performed after document retrieval. This does not match the contents of the document if any data conversion is performed and may result in application errors.

For implicit serialization, the XML declaration always matches the document contents.

8.1.3 Using file reference variables for efficient insert and retrieval

DB2 offers the use of file reference variables for XML data in the same manner as for LOBs. File reference variables are variables that point to documents in the file system and they allow you to refer to the documents without reading the contents of the documents into the memory of the application.

They can be used for efficient insert of XML documents that you receive from outside the application if you do not need to manipulate the contents before inserting, and for retrieval of documents that you want to pass on as they are without any modifications.

The variable declarations for file reference variables in COBOL can be seen in Example 8-5.

Example 8-5 XML file reference filterable in COBOL

```
01 DOC-AS-CLOB-FILE IS SQL TYPE IS XML AS CLOB-FILE.
01 DOC-AS-BLOB-FILE IS SQL TYPE IS XML AS BLOB-FILE.
01 DOC-AS-DBCLOB-FILE IS SQL TYPE IS XML AS DBCLOB-FILE.
```

These declarations are translated by the pre-compiler as shown in Example 8-6.

Example 8-6 XML file reference variables after transformation by the DB2 pre-compiler

```
01 DOC-AS-CLOB-FILE.
  49 DOC-AS-CLOB-FILE-NAME-LENGTH PIC S9(9) COMP-5 SYNC.
  49 DOC-AS-CLOB-FILE-DATA-LENGTH PIC S9(9) COMP-5.
  49 DOC-AS-CLOB-FILE-FILE-OPTION PIC S9(9) COMP-5.
  49 DOC-AS-CLOB-FILE-NAME PIC X(255).
01 DOC-AS-BLOB-FILE.
  49 DOC-AS-BLOB-FILE-NAME-LENGTH PIC S9(9) COMP-5 SYNC.
  49 DOC-AS-BLOB-FILE-DATA-LENGTH PIC S9(9) COMP-5.
  49 DOC-AS-BLOB-FILE-FILE-OPTION PIC S9(9) COMP-5.
  49 DOC-AS-BLOB-FILE-NAME PIC X(255).
01 DOC-AS-DBCLOB-FILE.
  49 DOC-AS-DBCLOB-FILE-NAME-LENGTH PIC S9(9) COMP-5 SYNC.
  49 DOC-AS-DBCLOB-FILE-DATA-LENGTH PIC S9(9) COMP-5.
  49 DOC-AS-DBCLOB-FILE-FILE-OPTION PIC S9(9) COMP-5.
  49 DOC-AS-DBCLOB-FILE-NAME PIC X(255).
```

To work with a file reference variable, you must initialize the fields concerning the file name, i.e. the name of the file and the length of the name, and the file option. The file option is used to signal which type of operation you want to perform on the file. These are supplied as constant declarations in COBOL, the possible values are shown in Table 8-2. The length of the file is provided by DB2 when writing to a file and does not have to be provided by the application.

Table 8-2 Options for file reference variables

FILE-OPTION	Constant	Operation
SQL-FILE-READ	2	Input from application to database. A regular file can be opened, read and closed
SQL-FILE-CREATE	4	Output from database to application. Creates a new file if one does not exist. If one does exist, an error is returned.

FILE-OPTION	Constant	Operation
SQL-FILE-APPEND	8	Output from database to application. Appends data to an existing file. If one does not exist, a new file is created.
SQL-FILE-OVERWRITE	16	Output from database to application. Overwrites an existing file. If one does not exist, a new file is created.

The initialization needed to read from a file named 'CAMT.STMT.XML' is shown in Example 8-7.

Example 8-7 Initialization of a file reference variable

```
MOVE 'CAMT.STMT.XML' TO DOC-AS-CLOB-FILE-NAME.
MOVE 13 TO DOC-AS-CLOB-FILE-NAME-LENGTH.
MOVE SQL-FILE-READ TO DOC-AS-CLOB-FILE-OPTION.
```

When the file reference variable is subsequently referred to, the specified file option is executed, thus if the file is used, such as in an insert operation, the file is read into DB2 and inserted without materializing in application memory.

Note: Apply PTF for currently open APAR PM25980 when using binary file reference and locator variables. It solves issues for non-UTF-8 encoded XML files.

8.2 The BankStmt application in COBOL

This section contains a step-by-step implementation of the BankStmt application in COBOL.

The purpose of the application is to demonstrate how to use COBOL with DB2 pureXML and we have chosen to emphasize the different steps and choices made that are related to XML, thus disregarding irrelevant code and components like presentation layer or end-user dialog. The application might therefore seem simple compared to a real-life application.

We shall use the same subset of the ISO20022 standard, namely the BankToCustomerStatement message, which has been used throughout the book.

The BankToCustomerStatement message is sent by the account servicer to an account owner or to a party authorized by the account owner to receive the message. It is used to inform the account owner, or authorized party, of the entries booked to the account, and to provide the owner with balance information on the account at a given point in time.

For the COBOL application, we have chosen the database schema that focuses on the statement as a business object. This implies that the BankToCustomerStatement is saved whole as one XML document, as opposed to shredding the message into smaller documents before saving them.

The level at which to define a document may be difficult to decide, but a good rule-of-thumb is that if the XML document matches a business object, it makes a good basis for the application because it contains the data that is typically used in each application component.

8.2.1 Setting up the environment

When creating a COBOL application, the initial work also includes setting up the environment. In this case we shall need tables for holding the data and the XML schema for validating the documents.

The BankToCustomerStatement schema

In order to understand the structure of the documents we work with, and to see the types of validation the schema imposes on this structure, we start by having a look at the schema for the BankToCustomerStatement.

The overall structure of a BankToCustomerStatement message as it looks when shown in a browser is shown in Figure 8-2. The elements can be expanded by clicking the '+' to the left or collapsed by clicking the '-'.

The outermost element is Document containing the element BkToCstmrStmt which in turn is made up from a GrpHdr element and a Stmt element. The Stmt element contains simple elements Id and CreDtTm, and complex elements FrToDt and Acct as well as a number of complex elements Bal and Ntry.

```

<?xml version="1.0" encoding="UTF-8" ?>
- <Document xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02">
- <BkToCstmrStmt>
  - <GrpHdr>
    <MsgId>AAAASESS-FP-STAT001</MsgId>
    <CreDtTm>2010-10-18T17:00:00+01:00</CreDtTm>
  </GrpHdr>
  - <Stmt>
    <Id>AAAASESS-FP-STAT001</Id>
    <CreDtTm>2010-10-18T17:00:00+01:00</CreDtTm>
    + <FrToDt>
    + <Acct>
    + <Bal>
    + <Bal>
    + <Ntry>
    + <Ntry>
    + <Ntry>
  </Stmt>
</BkToCstmrStmt>
</Document>

```

Figure 8-2 BankToCustomerStatement message as shown in a browser

The schema for messages of this format is available from the ISO20022 web site. It is quite extensive, and we shall not go into details, but let us have a look at a small subset of the schema that we shall be working with. Note, that even if you are not used to dealing with XML schemas, it is still possible to get an idea of what it says. The schema itself is also written in XML.

In Figure 8-3, we see the part of the schema that is concerned with the overall structure of the BankToCustomerStatement. We see the target namespace of the schema is *urn:iso:std:iso:20022:tech:xsd:camt.053.001.02*, and that the outermost element is named Document with type Document.

The type Document in turn is a complex type consisting of a sequence of elements with name BkToCstmrStmt of type BankToCustomerStatementV02 which is a complex type made up from another sequence of elements. This sequence contains an element of type GrpHdr and then one or more occurrences of the element Stmt.

```
- <xs:schema xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02"
  xmlns:xs="http://www.w3.org/2001/XMLSchema" elementFormDefault="qualified"
  targetNamespace="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02">
  <xs:element name="Document" type="Document" />
- <xs:complexType name="Document">
  - <xs:sequence>
    <xs:element name="BkToCstmrStmt" type="BankToCustomerStatementV02" />
  </xs:sequence>
</xs:complexType>
- <xs:complexType name="BankToCustomerStatementV02">
  - <xs:sequence>
    <xs:element name="GrpHdr" type="GroupHeader42" />
    <xs:element maxOccurs="unbounded" minOccurs="1" name="Stmt"
      type="AccountStatement2" />
  </xs:sequence>
</xs:complexType>
```

Figure 8-3 Subset of the BankToCustomerStatement schema

Schema registration

We need to register XML schemas in the DB2 XML schema repository to allow us to use DB2 to validate our XML documents.

The documents can either be validated explicitly using the DSN_XMLVALIDATE function when inserting or updating, or DB2 can perform the validation automatically on insert and update. The latter requires that the XML column is associated with an XML type modifier that holds information about which schemas to use for validation.

We have decided to use automatic schema validation for the following reasons:

- ▶ The application code is simpler.
- ▶ This ensures that all documents are validated, even though there might be several applications inserting or updating the XML documents.
- ▶ If the schema changes, we only need to alter the XML type modifier on the table, not all the insert and update applications.

The schema consists of a single document *camt.053.001.02.xsd* which we have obtained from the ISO20022 web site and have available on our workstation. To register the schema we have a number of options:

- ▶ Upload the schema to z/OS and register it through the command line processor (CLP) in Unix System Services (USS).
- ▶ Register the schema through the CLP in Windows on our workstation, connecting to DB2 on z/OS using DB2 Connect™.
- ▶ Use Optim Development Studio or similar development framework to register the schema. This also requires DB2 Connect to connect to DB2 for z/OS.

We choose to register the schema from the workstation using the CLP available in Windows. We first connect to DB2 using DB2 Connect, then register the schemas. The commands used are shown in Example 8-8.

For details on the other methods for schema registration, refer to 2.1.6, “XML schema repository and schema validation”

Example 8-8 CLP commands for registration of XML schema

```
db2 connect to DB0B user xmlr2 using passwd01
db2 register xmlschema 'camt.053.001.02.xsd' from 'camt.053.001.02.xsd' as
SYSXSR.camt_053_001_02
db2 complete xmlschema SYSXSR.camt_053_001_02
```

With the schema in place, we can create the tables needed for the application.

As stated earlier we have decided that the bank statement is what corresponds to a business object in this application, and we therefore need one table with one XML column, adding a few redundant columns for easy access and overview. The XML column is created with a type modifier referring to the XML schema that we just registered to allow for automatic schema validation.

The DDL for the tables is shown in Example 8-9. Note, that no indexes are created at this time. We add them when we have our programs and SQL queries in place to make sure that the best possible indexes are chosen.

Example 8-9 DDL for the table in the BankStmt application

```
CREATE TABLE BK_TO_CSTMTR_STMT
  (MSG_ID          VARCHAR(35) ,
   MSG_CRE_DT_TM   TIMESTAMP WITH TIME ZONE,
   BK_TO_CSTMTR_STMT XML
   (XMLSCHEMA ID SYSXSR.CAMT_053_001_02) NOT NULL) ;
```

8.2.2 Inserting XML documents

For the insert program we assume that we have the full BankToCustomerStatement message available in the file system. The insert program runs as a batch program, and the name of the file holding the XML document are passed as input in the JCL.

We also assume that the message is stored in EBCDIC, which matches the application code page, so the data is automatically converted correctly to UTF-8 if we use a single-byte character-based XML variable.

We shall need some data from the document to populate the redundant columns in the base table, but if we choose to populate those after the insert, we can do it in DB2 and we do not need to manipulate the contents of the XML file at all in the application program. This means that we can use a file reference variable to hold the XML.

The redundant columns are populated after the insert, so we only need to insert the XML document. We do, however, need to be able to identify the row after the insert in order to supply values for the redundant columns. To this end we shall use the unique column DB2_GENERATED_DOCID_FOR_XML COLUMN which is generated automatically by DB2 on insert of the XML document. We select this column from the final table of the insert statement.

The code needed for the initial insert is shown in Example 8-10.

Example 8-10 Insert a BankToCustomerStatement

```
01 WS-BLANK-STRING      PIC X(1) VALUE SPACES.
```

```

01 BK-STMT-DATA USAGE IS SQL TYPE IS XML AS CLOB-FILE.
01 DOC-ID          PIC S9(18) COMP-5.
...
ACCEPT BK-STMT-DATA-NAME FROM SYSIN.
INSPECT BK-STMT-DATA-NAME TALLYING BK-STMT-DATA-NAME-LENGTH
      FOR CHARACTERS BEFORE INITIAL WS-BLANK-STRING.
MOVE SQL-FILE-READ TO BK-STMT-DATA-FILE-OPTION.

EXEC SQL
  SELECT DB2_GENERATED_DOCID_FOR_XML INTO :DOC-ID
  FROM FINAL TABLE
  ( INSERT INTO BK_TO_CSTMR_STMT (BK_TO_CSTMR_STMT)
    VALUES (:BK-STMT-DATA) )
END-EXEC.

```

The XML data has now been validated and inserted, and all we need is to populate the two redundant columns. This is easily done by extracting the contents from the XML document using the XMLTABLE function.

The update, and thus the remainder of what is needed for the insert operation is shown in Example 8-11.

Example 8-11 Extracting key fields using XMLTABLE

```

EXEC SQL
  UPDATE BK_TO_CSTMR_STMT ST1
  SET ( MSG_ID , MSG_CRE_DT_TM ) =
  ( SELECT X.MSGID , X.CREDTM
  FROM BK_TO_CSTMR_STMT ST2,
  XMLTABLE (XMLNAMESPACES(DEFAULT
    'urn:iso:std:iso:2002:tech:xsd:camt.053.001.02',
    '/Document/BkToCstmrStmt/GrpHdr'
  PASSING ST2.BK_TO_CSTMR_STMT
  COLUMNS
  "MSGID"      VARCHAR(35) PATH 'MsgId',
  "CREDTM"    TIMESTAMP  PATH 'CreDtTm'
  ) X
  WHERE ST2.DB2_GENERATED_DOCID_FOR_XML = :DOC-ID )
  WHERE ST1.DB2_GENERATED_DOCID_FOR_XML = :DOC-ID
END-EXEC.

```

Note: XPath (as well as XQuery) expressions are case-sensitive.

To test the program we need two documents available in the file system: One, that is valid according to the XML schema and one that is not.

The JCL needed to run the program is shown in Example 8-12. The name of the file holding the XML document is given as an input parameter in the SYSIN DD card.

Example 8-12 JCL for running COBOL insert program

```

//PH02CS04 EXEC PGM=IKJEFT01,DYNAMNBR=20,COND=(4,LT)
//DBRMLIB DD DISP=SHR,DSN=DBOBM.DBRMLIB.DATA
//STEPLIB DD DISP=SHR,DSN=DBOBT.SDSNLOAD
//SYSTSPRT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*

```

```
//SYSUDUMP DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//REPORT DD SYSOUT=*
//SYSIN DD *
CAMT.BKSTMT.XML
//SYSTSIN DD *
DSN SYSTEM(DB0B)
  RUN PROGRAM(INSBKST) PLAN(COBXML) -
    LIBRARY('DBOBM.RUNLIB.LOAD')
  END
//CARDIN DD *
```

We first run it with the XML document that is not valid. To produce the non-valid document we have omitted the GrpHdr element altogether; the XML schema states that this is a required element as we saw in the previous section.

Running the program with this document gives us the error shown in Example 8-13.

Example 8-13 Validation error on insert

```
DSNT408I SQLCODE = -20399, ERROR: ERROR ENCOUNTERED DURING XML VALIDATION:
LOCATION 184; TEXT An expected element match was not found.RC=0018,RSN=8604;
XSRID 144.
DSNT418I SQLSTATE = 2201R SQLSTATE RETURN CODE
DSNT415I SQLERRP = DSNNOPAR SQL PROCEDURE DETECTING ERROR
DSNT416I SQLERRD = -510 0 0 -1 0 0 SQL DIAGNOSTIC INFORMATION
DSNT416I SQLERRD = X'FFFFFFE02' X'00000000' X'00000000' X'FFFFFFF' X'00
INFORMATION
```

We see that the error is detected in location 184 of the document, and that an expected element match was not found there. By inspecting the document we find that this is the first location after the start tag of the BkToCstmrStmt, and this is exactly the place where the GrpHdr should be according to the schema.

We then run it with the valid document and the program completes as expected. To verify that the document has been validated, we use the DB2-supplied SQL function XMLXSROBJECTID which takes an XML column and returns the XSR object identifier that was used to validate the XML document - or 0 if the document has not been validated. The identifier can then be looked up in the SYSXSR.SYSOBJECTS catalog table as shown in Example 8-14.

Example 8-14 Determining whether an XML document has been validated

```
-----+-----+-----+-----+-----+-----+-----+
SELECT B.ID, S.XSROBJECTNAME
FROM BK_TO_CSTMRTMT B
, SYSIBM.XSROBJECTS S
WHERE S.XSROBJECTID = XMLXSROBJECTID(B.BK_TO_CSTMRTMT)
-----+-----+-----+-----+-----+-----+-----+
          ID XSROBJECTNAME
-----+-----+-----+-----+-----+-----+-----+
          5 CAMT_053_001_02
DSNE610I NUMBER OF ROWS DISPLAYED IS 1
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100
-----+-----+-----+-----+-----+-----+-----+
```

8.2.3 Updating XML documents

Now let us look at a small program that updates a BankToCustomerStatement.

A BankToCustomerStatement can normally be sent either to the owner of the account, or to another recipient. This is modelled by having a MsgRcpt element in the GrpHdr of the document which is filled in only if the recipient is different from the owner. An example of a recipient is shown in Figure 8-4.

```
- <MsgRcpt>
  <Nm>Penelope Keith</Nm>
- <PstAdr>
  <StrtNm>Bailey Avenue</StrtNm>
  <BldgNb>555</BldgNb>
  <PstCd>95141</PstCd>
  <TwnNm>San Jose</TwnNm>
</PstAdr>
</MsgRcpt>
```

Figure 8-4 Message recipient of a BankToCustomerStatement

Now imagine that we want to change the recipient of a BankToCustomerStatement. For this purpose we use the DB2 function XMLMODIFY which can update part of an XML document.

The XMLMODIFY function has three subfunctions. It can

- ▶ Insert data into an XML document
- ▶ Replace data in an XML document
- ▶ Delete data from an XML document

Depending on the situation we might want to use any one of these functions.

- ▶ If the recipient is the owner and we want to change it to someone else, we should insert a MsgRcpt element
- ▶ If the recipient is not the owner and we want to change to someone else who is not the owner, we should replace the MsgRcpt element
- ▶ If the recipient is someone other than the owner and we want to change it to the owner, we should delete the MsgRcpt element

Note: The use of the XMLMODIFY function to update parts of an XML document is supported for tables with the multi-versioning format introduced in DB2 10 only.

A table has the multi-versioning format if it is created in DB2 10 NFM, has an XML column, and resides in a universal table space. Or if it is created in DB2 9, resides in a universal table space and all the XML columns are added to the table in DB2 10 NFM.

For now, let us assume that the MsgRcpt element is once again available in a file which we can access through a file reference variable. We input to the program the name of this file, an ID of the XML message we want to change, and a choice of function to perform - replace, insert or update.

We shall also consider an alternative to this approach, namely to input the raw data and then build an XML element using COBOL features. See 8.3.1, "Generation of XML documents in COBOL" on page 176.

The updating program is shown in Example 8-15. The program inputs the MSG_ID of the row to update, a number indicating which function to perform and the name of a file holding a MsgRcpt element. Depending on the function choice it executes one of three SQL statements using XMLMODIFY to update the XML document. It either inserts the MsgRcpt element, replaces an existing MsgRcpt element, or deletes an existing MsgRcpt element.

Example 8-15 COBOL program for updating a BkToCstmrStmt with a new MsgRcpt

```

WORKING-STORAGE SECTION.
01 REPLACE-MSG-RCPT   PIC 9 VALUE 1.
01 INSERT-MSG-RCPT   PIC 9 VALUE 2.
01 DELETE-MSG-RCPT   PIC 9 VALUE 3.
01 FUNCTION-CHOICE    PIC 9.
    EXEC SQL INCLUDE SQLDA END-EXEC.
    EXEC SQL INCLUDE SQLCA END-EXEC.
01 WS-BLANK-STRING    PIC X(1) VALUE SPACES.
01 NEW-RCPT USAGE IS SQL TYPE IS XML AS CLOB-FILE.
01 MSGID              PIC X(35).
...
MAIN SECTION.
    EXEC SQL WHENEVER SQLERROR   GOTO DBERROR END-EXEC.
    EXEC SQL WHENEVER SQLWARNING GOTO DBERROR END-EXEC.

    ACCEPT MSGID FROM SYSIN.
    ACCEPT FUNCTION-CHOICE FROM SYSIN.
    ACCEPT NEW-RCPT-NAME FROM SYSIN.
    INSPECT NEW-RCPT-NAME TALLYING NEW-RCPT-NAME-LENGTH
        FOR CHARACTERS BEFORE INITIAL WS-BLANK-STRING.
    MOVE SQL-FILE-READ TO NEW-RCPT-FILE-OPTION.

    EVALUATE FUNCTION-CHOICE
        WHEN REPLACE-MSG-RCPT PERFORM REPLACE-RCPT
        WHEN INSERT-MSG-RCPT PERFORM INSERT-RCPT
        WHEN DELETE-MSG-RCPT PERFORM DELETE-RCPT
        WHEN OTHER DISPLAY "OTHER" FUNCTION-CHOICE
    END-EVALUATE.

REPLACE-RCPT.
    EXEC SQL
        UPDATE BK_TO_CSTMR_STMT
        SET BK_TO_CSTMR_STMT =
        XMLMODIFY (
            'declare default element namespace
-         '"urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
-         'replace node
-         '/Document/BkToCstmrStmt/GrpHdr/MsgRcpt
-         'with $rcp/NewRcpt/MsgRcpt',
            :NEW-RCPT AS "rcp"
        )
        WHERE MSG_ID = :MSGID
    END-EXEC.

INSERT-RCPT.
    EXEC SQL
        UPDATE BK_TO_CSTMR_STMT
        SET BK_TO_CSTMR_STMT =

```



```

XMLMODIFY (
  'declare default element namespace
-  'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
-  'insert node $rcp/NewRcpt/MsgRcpt after
-  '/Document/BkToCstmrStmt/GrpHdr/CreDtTm',
  :NEW-RCPT AS "rcp"
)
WHERE MSG_ID = :MSGID
END-EXEC.

DELETE-RCPT.
EXEC SQL
  UPDATE BK_TO_CSTMTR_STMT
  SET BK_TO_CSTMTR_STMT =
  XMLMODIFY (
    'declare default element namespace
-  'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
-  'delete node
-  '/Document/BkToCstmrStmt/GrpHdr/MsgRcpt'
  )
  WHERE MSG_ID = :MSGID
END-EXEC.

```

Make a note of the following two points:

- ▶ For both the replace and insert expressions, a path of \$rcp/Newrcpt/MsgRcpt is used instead of just \$rcp to indicate the new element. This is because we have wrapped the XML element in a container element, in this case NewRcpt, which is not to be inserted into DB2. This is necessary when inserting more than one element and therefore a good practice to adapt in general for consistency.
- ▶ XPath expressions may be quite long, spanning several lines, especially if many namespace declarations are needed. The XPath expressions in the program are enclosed in apostrophes, and an apostrophe is used together with the continuation character '-' to indicate that the expression continues on the next line. This requires the program to be precompiled with the APOST and APOSTSQL options.

The program was tested with function 1, for insert of a new MsgRcpt element, and a file containing the MsgRcpt element shown in Figure 8-4 on page 167. This produced the error shown in Example 8-16.

Example 8-16 SQL error when updating XML document with MsgRcpt element

```

DSNT408I  SQLCODE = -20399, ERROR:  ERROR ENCOUNTERED DURING XML VALIDATION:
        LOCATION 237; TEXT An element is not in the choice.RC=0018,RSN=8608;
        XSRID 144.
DSNT418I  SQLSTATE  = 2201R SQLSTATE RETURN CODE
DSNT415I  SQLERRP   = DSNNOPAR SQL PROCEDURE DETECTING ERROR
DSNT416I  SQLERRD   = -510 0 0 -1 0 0 SQL DIAGNOSTIC INFORMATION
DSNT416I  SQLERRD   = X'FFFFFFE0' X'00000000' X'00000000' X'FFFFFFF'
        X'00000000' X'00000000' SQL DIAGNOSTIC INFORMATION

```

The error states that an element that begins in location 237 of the XML document is not allowed in that place. This turns out to be the MsgRcpt element which is not surprising as this is the only place where we have changed the document.

The `MsgRcpt` element is inserted after the `CreDtTm` element in the `GrpHdr` element. Let us compare this to the schema definition of the `GrpHdr` element which is shown in Figure 8-5. In fact, this is a definition of the complex type `GroupHeader42` which describes the `GrpHdr` element. We see that it comprises a sequence with a `MsgId` element, followed by a `CreDtTm` element, which again is followed by an optional `MsgRcpt` element and two other optional elements.

```
- <xs:complexType name="GroupHeader42">
  - <xs:sequence>
    <xs:element name="MsgId" type="Max35Text" />
    <xs:element name="CreDtTm" type="ISODateTime" />
    <xs:element maxOccurs="1" minOccurs="0" name="MsgRcpt" type="PartyIdentification32" />
    <xs:element maxOccurs="1" minOccurs="0" name="MsgPgntn" type="Pagination" />
    <xs:element maxOccurs="1" minOccurs="0" name="AddtlInf" type="Max500Text" />
  </xs:sequence>
```

Figure 8-5 Schema definition for the `GrpHdr` element

In other words, the location of the `MsgRcpt` element is not in conflict with the schema definition, so why does DB2 not recognize the element name as valid in the context?

The explanation is related to namespaces. We recall that namespaces are a mechanism for ensuring uniqueness of element names, so that two elements in different domains and possibly with different structure and contents are not confused even though they have the same element name. By associating with each element a namespace, we guarantee that we know which one we are referencing. This association can be done either explicitly with a prefix to the element name, or implicitly by declaring a default element namespace.

The `BankToCustomerStatement` has the default namespace of `urn:iso:std:iso:20022:tech:xsd:camt.053.001.02`, and therefore all the elements belonging to this message have the same namespace unless another namespace is explicitly given. The `MsgRcpt` element we attempted to insert into the document had no namespaces associated with it, and therefore it is not the same element as the one in the `BankToCustomerStatement` schema. Hence, it is not valid in the context.

We alter the contents of the file containing the `MsgRcpt` element as shown in Figure 8-6 and attempt to run the update again. This time the namespace declaration matches the default element declaration, and the insert succeeds.

```
- <NewRcpt xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02">
  - <MsgRcpt>
    <Nm>Penelope Keith</Nm>
    - <PstAdr>
      <StrtNm>Bailey Avenue</StrtNm>
      <BldgNb>555</BldgNb>
      <PstCd>95141</PstCd>
      <TwnNm>San Jose</TwnNm>
    </PstAdr>
  </MsgRcpt>
</NewRcpt>
```

Figure 8-6 `MsgRcpt` element with namespace declaration

8.2.4 Querying XML documents

For retrieving XML from DB2, there are two options. Either the data is retrieved as XML, or else the data is converted to simple SQL types either as a result of a cast operation, or from using the XMLTABLE function.

In Example 8-17 it is shown how to select a whole XML document with a given ID from the BK_TO_CSTMRT_STMT table and write it to a file using a file reference variable.

Example 8-17 Retrieval of an XML document to a file

```

01 WS-BLANK-STRING    PIC X(1) VALUE SPACES.
01 MSGID              PIC X(35).
01 BK-STMT-FILE USAGE IS SQL TYPE IS XML AS CLOB-FILE.
   EXEC SQL INCLUDE SQLDA END-EXEC.
   EXEC SQL INCLUDE SQLCA END-EXEC.
...
MAIN SECTION.
   EXEC SQL WHENEVER SQLERROR  GOTO DBERROR END-EXEC.
   EXEC SQL WHENEVER SQLWARNING GOTO DBERROR END-EXEC.

   ACCEPT MSGID FROM SYSIN.

   ACCEPT BK-STMT-FILE-NAME FROM SYSIN.
   INSPECT BK-STMT-FILE-NAME TALLYING BK-STMT-FILE-NAME-LENGTH
     FOR CHARACTERS BEFORE INITIAL WS-BLANK-STRING.
   MOVE SQL-FILE-OVERWRITE TO BK-STMT-FILE-FILE-OPTION.

   EXEC SQL
     SELECT BK_TO_CSTMRT_STMT INTO :BK-STMT-FILE
     FROM BK_TO_CSTMRT_STMT
     WHERE MSG_ID = :MSGID
   END-EXEC.

```

Now imagine that we want to make a list of all entries created after a given date. The list is to contain the statement id, amount, currency, credit-debit indicator and datetime.

We want to extract these values into simple SQL host variables of type DECIMAL, CHAR and TIMESTAMP. In this case it is actually transparent to the COBOL application that we are working with XML data, because all the XML manipulation takes place in DB2 via EXEC SQL statements. Even for the host variable declarations, we do not have to consider the different XML alternatives.

See Example 8-18 for a program that extracts the entries for all bank statement and places them in relational host variables. The program inputs a timestamp and selects data from entries that have a timestamp later than this timestamp. This is done by passing the value of the timestamp as a parameter to the XPath expression, and then using it in the predicate where it is compared to the timestamp of each entry.

It uses XMLTABLE to get a relational view of each entry, of which there may be several per BankToCustomerStatement, so potentially there are more rows returned than rows in the table. These are then filtered by the predicate, so potentially there could also be fewer rows than rows in the table. The values of these are then placed in relational host variables for further processing.

Example 8-18 Retrieval of data in relational format from an XML document

```

01 FROM-TIME          PIC X(19).
01 STMT-ID            PIC X(20).
01 NTRY-TIME         PIC X(26).
01 AMOUNT            PIC S9(9)V9(2) COMP-3.
01 CREDIT-DEBIT     PIC X(4).
01 CURRENCY-NM      PIC X(4).
...
MAIN SECTION.
EXEC SQL WHENEVER SQLERROR  GOTO DBERROR END-EXEC.
EXEC SQL WHENEVER SQLWARNING GOTO DBERROR END-EXEC.

ACCEPT FROM-TIME FROM SYSIN.

EXEC SQL
  DECLARE C1 CURSOR FOR
  SELECT X.STMT, X.DTTM, X.CCY , X.AMT , X.CREDBTIND
  FROM BK_TO_CSTMRT STMT S,
  XMLTABLE (XMLNAMESPACES(DEFAULT
  'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02',
  '/Document/BkToCstmrStmnt/Stmnt/Ntry[BookgDt/DtTm>$tm]')
  PASSING S.BK_TO_CSTMRT, TIMESTAMP(:FROM-TIME) AS "tm"
  COLUMNS
  "STMT"      CHAR(20)      PATH './Id',
  "DTTM"     CHAR(26)      PATH 'BookgDt/DtTm',
  "CCY"      CHAR(4)       PATH 'Amt/@Ccy',
  "AMT"      DECIMAL(11,2) PATH 'Amt',
  "CREDBTIND" CHAR(4)     PATH 'CdtDbtInd'
  ) X
  END-EXEC.
EXEC SQL
  OPEN C1
  END-EXEC.
EXEC SQL
  FETCH C1 INTO :STMT-ID ,
                :NTRY-TIME ,
                :CURRENCY-NM ,
                :AMOUNT ,
                :CREDIT-DEBIT
  END-EXEC.
PERFORM WRITE-AND-FETCH
  UNTIL SQLCODE IS NOT EQUAL TO ZERO.
EXEC SQL WHENEVER NOT FOUND GOTO CLOSEC1  END-EXEC.
CLOSEC1.
EXEC SQL CLOSE C1  END-EXEC.

```

8.2.5 Designing indexes

We recall that XML indexes can be utilized by XMLEXISTS and XMLTABLE functions, so only XPath patterns used in one of these functions are candidates for index use.

In the COBOL application we have not used the XMLEXISTS function at all. The XMLTABLE function has been used twice, but only the one extracting entries from the bank statement contained a predicate.

The pattern used for the context node in this expression is the candidate for index access. It is shown in Example 8-19.

Example 8-19 Candidate index pattern for the BankStmt application

```
'/Document/BkToCstmrStmt/Stmt/Ntry[BookgDt/DtTm>$tm]'
```

We need to include the whole path down to and including the DtTm element as this is the one the predicate evaluates, and we need to include namespace declarations in the index. The resulting index can be seen in Example 8-20.

Example 8-20 XML index on DtTm elements

```
CREATE INDEX IXMLNTRY
ON BK_TO_CSTMTR_STMT(BK_TO_CSTMTR_STMT)
GENERATE KEY USING XMLPATTERN
'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
/Document/BkToCstmrStmt/Stmt/Ntry/BookgDt/DtTm'
AS SQL TIMESTAMP
```

Note: The support for date and time data types and functions in XML functions, including TIME, DATE and TIMESTAMP as data types for indexes, requires DB2 10 NFM.

To ensure that the index is utilized by the COBOL program, we run the Runstats utility on the table and XML index, and then do an explain of the program. In Example 8-21 a select of a few essential columns from the plan table is shown for the program after and before the creation of the index.

The access type for the access to table BK_TO_CSTMTR_STMT is DX which is an indication that an XML index is being used for access, the access name is IXMLNTRY which is the index we just created.

Example 8-21 Access path using the index IXMLNTRY

```
SELECT CREATOR, TNAME, METHOD, ACESSTYPE, ACCESSCREATOR, ACCESSNAME
FROM PLAN_TABLE
WHERE PROGNAME = 'GETNTRY'
;
```

CREATOR	TNAME	METHOD	ACESSTYPE	ACCESSCREATOR	ACCESSNAME
XMLR2	BK_TO_CSTMTR_ST	0	DX	XMLR2	IXMLNTRY
XMLR2	X	1	R		

```
DSNE610I NUMBER OF ROWS DISPLAYED IS 2
DSNE612I DATA FOR COLUMN HEADER TNAME COLUMN NUMBER 2 WAS TRUNCATED
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100
```

8.2.6 Schema evolution

As part of the application life-cycle it is expected that an XML format will evolve over time. In this section, we show an example of a schema change, identify which components are affected, and demonstrate how to change the application accordingly.

We recall that a BankToCustomerStatement can be sent either to the owner of the account, or to another recipient and that this is modelled by having a MsgRcpt element in the GrpHdr of the document which is filled in only if the recipient is different from the owner.

Imagine now that you wanted to send the same message to more than one recipient, so in the XML document you would have more than one MsgRcpt. What would be the impact on the schema and your applications? The schema definition of the GrpHdr is shown in Figure 8-7.

```
- <xs:complexType name="GroupHeader42">
  - <xs:sequence>
    <xs:element name="MsgId" type="Max35Text" />
    <xs:element name="CreDtTm" type="ISODatetime" />
    <xs:element maxOccurs="1" minOccurs="0" name="MsgRcpt" type="PartyIdentification32" />
    <xs:element maxOccurs="1" minOccurs="0" name="MsgPgntn" type="Pagination" />
    <xs:element maxOccurs="1" minOccurs="0" name="AddtlInf" type="Max500Text" />
  </xs:sequence>
</xs:complexType>
```

Figure 8-7 Schema definition of GrpHdr element

We notice that the current schema definition does not allow more than occurrence of the MsgRcpt element, this is determined by the maxOccurs="1" attribute, and if you tried to insert a document containing e.g. two message recipients, you would get a validation error. So the first thing that has to change is the schema definition.

To alter the schema definition to allow for multiple occurrences of the MsgRcpt element, all you need to do is change the maxOccurs="1" clause of the element to maxOccurs="unbounded". The result is shown in Figure 8-8.

```
- <xs:complexType name="GroupHeader42">
  - <xs:sequence>
    <xs:element name="MsgId" type="Max35Text" />
    <xs:element name="CreDtTm" type="ISODatetime" />
    <xs:element maxOccurs="unbounded" minOccurs="0" name="MsgRcpt"
      type="PartyIdentification32" />
    <xs:element maxOccurs="1" minOccurs="0" name="MsgPgntn" type="Pagination" />
    <xs:element maxOccurs="1" minOccurs="0" name="AddtlInf" type="Max500Text" />
  </xs:sequence>
</xs:complexType>
```

Figure 8-8 Revised schema definition for GrpHdr with multiple MsgRcpt elements

We give the schema another version, e.g. SYSXSR.CAMT_053_001_03 and register it to DB2 in the same way as we did the original schema. See Example 8-8 on page 164 for details.

We then need to alter the table definition so that the version 3 schema is used for automatic validation. We do not want to remove the original schema because the data already in the table was validated against this schema, and removing it would remove the audit trail. It would also leave the table in check pending because the rows would not have been validated against a schema mentioned in the XML type modifier. Instead we add the new version to the type modifier of the XML column on top of the original schema.

When updating or inserting new XML documents, these are automatically validated against the latest version of the type modifier when using automatic validation.

The DDL needed to extend the type modifier of the XML column is shown in Example 8-22.

Example 8-22 Adding a new schema to an XML type modifier

```
ALTER TABLE BK_TO_CSTMRTMT
ALTER DOCUMENT
SET DATA TYPE XML
(XMLSCHEMA ID SYSXSR.CAMT_053_001_02,
ID SYSXSR.CAMT_053_001_03)
```

The table and schema repository are now set to handle the updated schema.

Application changes

On the application side, the changes needed are of course influenced by the degree to which the fields involved in the change are used. The following examples nevertheless show that XML in many ways are quite robust to minor changes.

The insert application takes a whole XML document from the file system and inserts it into DB2 without manipulating it at all. The data validation is performed by DB2, but as we have altered the XML type modifier to include the new schema, this is already taken care of. Hence, no changes are necessary in the insert application.

The update application alters the message recipient of the statement. Instead of altering it, we might in future want to just add another message recipient as this is now allowed by the new schema. The COBOL program already contains the option of inserting a new message recipient, let us investigate it to see whether it may be used for inserting additional message recipients next to existing ones.

The XMLMODIFY expression used for insert is shown in Example 8-23. Note, that it inserts the `MsgRcpt` element right after the `CreDtTm` element. This means that if one or more `MsgRcpt` elements are already there, the new element is placed before these. This is valid according to the schema, so unless any significance is given to the order of the `MsgRcpt` elements, the application needs no changes to cater for the schema change.

In fact, because we chose to wrap the `MsgRcpt` element in a container element called `NewRcpt`, we are now also able to use the insert function to insert several new message recipients at one time, if we found that desirable.

Example 8-23 Insert a MsgRcpt element after the CreDtTm element

```
EXEC SQL
  UPDATE BK_TO_CSTMRTMT
  SET BK_TO_CSTMRTMT =
  XMLMODIFY (
    'declare default element namespace
-   'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
-   'insert node $rcp/NewRcpt/MsgRcpt after
-   '/Document/BkToCstmrtmt/GrpHdr/CreDtTm',
    :NEW-RCPT AS "rcp"
  )
  WHERE ID = :UPD-ID
END-EXEC.
```

None of the query applications make use of the message recipients, so no changes are necessary here.

Comparing to relational schema change

In summary, the only changes needed in the COBOL BankStmt application to allow multiple message recipients were the following changes in the environment:

- ▶ In the schema definition, change the attribute of the Msgrcpt element from maxOccurs =“1” to maxOccurs =“unbounded”
- ▶ Register the resulting schema in DB2 with a new version number
- ▶ Extend the type modifier of the XML column to include the new schema version

All of these changes can be performed as online changes.

No changes at all were necessary on the application side.

If the database schema had been relational, the required changes would likely be much more extensive involving a new table for the message recipient data as well as converting existing data to make them available in the new table.

This, in turn, would require application changes, allowing the application to insert and update data in the new table.

8.3 COBOL functions for manipulating XML

COBOL also offers support for XML, independently of the pureXML support in DB2.

This functionality may in some cases complement the DB2 functionality when basing the database design on pureXML.

We shall give a brief introduction to the most important concepts, namely parsing, generation and validation of XML documents in native COBOL. For more detailed information, refer to *Enterprise COBOL for z/OS Version 4.2 Programming Guide*, SC23-8529-01.

8.3.1 Generation of XML documents in COBOL

The DB2 pureXML functionality for creating XML documents from relational data is offered by the various publishing functions like XMLDOCUMENT, XMLELEMENT, XMLNAMESPACE etc. These functions were studied in 2.1.2, “SQL/XML language” on page 24.

Correspondingly, COBOL offers support for generation of XML documents from COBOL structures through the XML GENERATE statement.

This function takes as input a data item which is typically a group, and generates as output an XML document with similar structure as the input data item. The resulting element names are taken from the names in the group data item, and the resulting element contents is taken from the contents of these variables.

In Example 8-24 we see how to use XML GENERATE to generate a MsgRcpt element from a variable structure.

Example 8-24 COBOL program for generation of the MsgRcpt element

```
01 NewRcpt.
   02 MsgRcpt.
       05 Nm          PIC X(20) Value 'Pamela Woods'.
       05 Pst1Adr.
```



```

        10 StrtNm PIC X(20) Value '555'.
        10 BldgNb PIC X(20) Value 'Bailey Avenue'.
        10 PstCd PIC X(20) Value '95141'.
        10 TwNm PIC X(20) Value 'San Jose'.
01 NS PIC X(50)
    VALUE 'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02'.
01 NEW-RCPT USAGE IS SQL TYPE IS XML AS CLOB(1K).
    EXEC SQL DECLARE :NEW-RCPT VARIABLE CCSID 1208 END-EXEC.
...
XML GENERATE NEW-RCPT-DATA FROM NewRcpt
COUNT IN NEW-RCPT-LENGTH
WITH ENCODING 1208
NAMESPACE IS NS
END-XML.

```

There are a couple of things we should note from this example:

- ▶ The result of the XML GENERATE operation is stored in the variable NEW-RCPT defined as CLOB AS XML. To initialize the length of this variable the keyword COUNT IN is used.
- ▶ The XML document is generated with code page UTF-8 to avoid code page conversion. This is done by use of the keyword WITH ENCODING in the XMLPARSE statement. To pass this information to DB2, an explicit CCSID declaration of the host variable:NEW-RCPT is made.
- ▶ The namespace is provided in a separate alphanumeric variable and included in the generation with the keyword NAMESPACE IS. This is optional for the XML GENERATE statement.
- ▶ The element names are copied exactly as declared in the COBOL program so if mixed case element names are needed, make sure that the variable declarations are in mixed case.

The result of the generation is shown in Figure 8-9.

```

- <NewRcpt xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02">
- <MsgRcpt>
  <Nm>Penelope Keith</Nm>
  - <PstAdr>
    <StrtNm>Bailey Avenue</StrtNm>
    <BldgNb>555</BldgNb>
    <PstCd>95141</PstCd>
    <TwNm>San Jose</TwNm>
  </PstAdr>
</MsgRcpt>
</NewRcpt>

```

Figure 8-9 *MsgRcpt* element created with XML GENERATE

This provides us with an alternative to the program for updating the *MsgRcpt* of a bank statement which was discussed in 8.2.3, “Updating XML documents” on page 167. Instead of assuming that the new *MsgRcpt* is provided as an XML element in a text file, we can input the text values and build the XML element ourselves.

The XML GENERATE function also has the option of generating the values as attributes instead of text elements by using the keyword WITH ATTRIBUTES. This would result in the XML element shown in Figure 8-10.

```

- <NewRcpt xmlns="urn:iso:std:iso:2002:tech:xsd:camt.053.001.02">
- <MsgRcpt Nm="Pamela Woods">
  <PstlAdr StrtNm="Bailey Avenue" BldgNb="555" PstCd="95141" TwnNm="San Jose" />
</MsgRcpt>
</NewRcpt>

```

Figure 8-10 *MsgRcpt* element created with XML GENERATE WITH ATTRIBUTES

In most cases, the simple XML GENERATE is probably the better choice. At the moment, there is no support for generating XML documents with both text elements and attributes.

8.3.2 Shredding XML documents in COBOL

DB2 offers support for shredding XML documents into relational data through the XMLTABLE function. This function assumes that the XML document is stored in DB2.

COBOL has similar support for XML documents stored in COBOL variables, this is offered through the XML PARSE statement.

This statement takes as input an XML document and a parsing procedure that handles the events that occur during parsing. This enables you to shred the document into COBOL variables, e.g. an alphanumeric group with the same structure as the XML document or separate elementary data items. In addition, you could just process the data directly without saving the XML contents into variables.

The parsing procedure has to be written in the application program using the different XML events provided by the XML parser as it goes through the document.

In Example 8-25, we see how to use XML PARSE with a processing procedure to parse an *MsgRcpt* element into a variable structure.

Example 8-25 *COBOL program for shredding a MsgRcpt element into variables*

```

01 MsgRcpt.
  05 Nm PIC X(20).
  05 PstlAdr.
    10 StrtNm PIC X(20).
    10 BldgNb PIC X(20).
    10 PstCd PIC X(20).
    10 TwnNm PIC X(20).
01 NS          PIC X(50).
01 RCPT       PIC X(207).
01 CURRENT-ELEMENT PIC X(40).
...
XML PARSE NEW-RCPT
  PROCESSING PROCEDURE GET-DATA
END-XML.
...
GET-DATA.
  EVALUATE XML-EVENT
    When 'START-OF-ELEMENT'
      Move XML-Text to current-element
    When 'CONTENT-CHARACTERS'
      EVALUATE current-element
        When 'Nm'

```

```

        Move XML-TEXT TO Nm
    When 'Pst1Adr'
        Move XML-TEXT TO Pst1Adr
    When 'StrtNm'
        Move XML-TEXT TO StrtNm
    When 'BldgNb'
        Move XML-TEXT TO BldgNb
    When 'PstCd'
        Move XML-TEXT TO PstCd
    When 'TwnNm'
        Move XML-TEXT TO TwnNm
    When Other
        Continue
    End-evaluate
When 'ATTRIBUTE-NAME'
    Continue
When 'ATTRIBUTE-CHARACTERS'
    Continue
When 'EXCEPTION'
    DISPLAY 'Exception code: ' XML-CODE
END-EVALUATE.

```

The XML PARSE statement was introduced in COBOL before the pureXML format was available in DB2, and is useful for shredding XML documents saved as LOBs in DB2.

In general, we recommend to save XML data as pureXML especially if you need to query the contents of those data, and in that case the shredding is more readily done by the XMLTABLE function in DB2.

8.3.3 Validation of XML documents in COBOL

Finally, COBOL also offers validation of XML documents against an XML schema through a variant of the XML PARSE statement. This requires Enterprise COBOL for z/OS V4.2.

The schema does not have to be registered anywhere, but it does have to be in a pre-processed format known as *Optimized Schema Representation (OSR)*. This can be done from Unix System Services with a command like the one in Example 8-26. First, the schema is copied to USS from TSO, next the OSR document is generated.

Example 8-26 Converting a schema to OSR format

```

cp -B "'/'XMLR2.BKSTMT.XSD'" /u/xmlr2/bkstmt.xsd
xsdosrg -v -o /u/xmlr2/bkstmt.osr /u/xmlr2/bkstmt.xsd

```

We have extended the XML PARSE statement in Example 8-25 on page 178 with the validating phrase as shown in Example 8-27.

Example 8-27 XMLPARSE with schema validation

```

CONFIGURATION SECTION.
SPECIAL-NAMES.
    XML-SCHEMA RSCHEMA IS 'DDSCHEMA'.
...
XML PARSE RCPT
    WITH ENCODING 1208
    VALIDATING WITH FILE RSCHEMA

```

```
PROCESSING PROCEDURE GET-DATA  
END-XML.
```

The schema declaration in the SPECIAL-NAMES section allows us to associate the schema name RSCHEMA with an external file containing the schema. This can then be supplied in as a DD statement in the JCL to run the COBOL program as shown in Example 8-28.

Example 8-28 DD statement for supplying a schema to the COBOL program

```
//G0.DDSHEMA DD PATH='/u/xmlr2/bkstmt.osr'
```

As mentioned earlier, the automatic schema validation that can be set up in DB2 just by associating an XML type identifier with the XML column, is very robust to schema changes and very easy to work with. This makes it generally a better choice than explicit schema validation in both DB2 and COBOL.



Utilities with XML

In this chapter we introduce the use of DB2 utilities with XML data types. The utilities handle XML objects in a way similar to the way they handle LOB objects, but for some utilities you need to specify certain XML keywords.

This chapter contains the following sections:

- ▶ CHECK DATA
- ▶ CHECK INDEX
- ▶ For UNLOAD, SPANNED with BINARY XML performs significantly better than other formats. Use SPANNED if possible rather than file references.
- ▶ COPY
- ▶ COPYTOCOPY
- ▶ EXEC SQL
- ▶ LISTDEF
- ▶ LOAD
- ▶ For UNLOAD, SPANNED with BINARY XML performs significantly better than other formats. Use SPANNED if possible rather than file references.
- ▶ QUIESCE
- ▶ REBUILD INDEX
- ▶ RECOVER INDEX and RECOVER TABLESPACE
- ▶ REORG INDEX and REORG TABLESPACE
- ▶ REPAIR
- ▶ REPORT
- ▶ RUNSTATS
- ▶ UNLOAD
- ▶ DSNTIAUL
- ▶ DSN1COPY

We discuss only those features that are directly related to XML.

9.1 CHECK DATA

We discussed in Example 5-3 on page 76 and Example 5-4 on page 77 when an XML table space can be placed in CHECK-pending state. You should run the CHECK DATA utility to reset the CHECK-pending state.

The CHECK DATA utility checks XML relationships and can check the consistency between a base table space and the corresponding XML table spaces. If the base table space is not consistent with any related XML table spaces, CHECK DATA reports the error.

The default behavior of CHECK DATA is to check all objects that are in CHECK-pending status (SCOPE PENDING). However, you can limit the scope of checking by specifying SCOPE REFONLY to check only the base tables or SCOPE AUXONLY to check XML and LOB objects.

You can specify the action that DB2 performs when it finds an error in XML columns by using keyword XMLERROR. XMLERROR REPORT reports the error only and XMLERROR INVALIDATE reports the error and sets the column in error to an invalid status.

You can specify the action that DB2 performs when it finds an error in LOB or XML columns by using keyword AUXERROR. AUXERROR REPORT reports the error only and AUXERROR INVALIDATE reports the error and sets the column in error to an invalid status.

You do not normally specify both XMLERROR and AUXERROR.

CHECK DATA utility has the following features to support XML data.

- ▶ Check consistency between the base table space and the NODEID index.
- ▶ Check consistency between the XML table space and the NODEID index .
- ▶ Check consistency in the document structure for each XML document.
- ▶ Validate schema if XML columns have a type modifier.

Figure 9-1 shows the CHECK DATA utility syntax for keywords introduced in DB2 10.

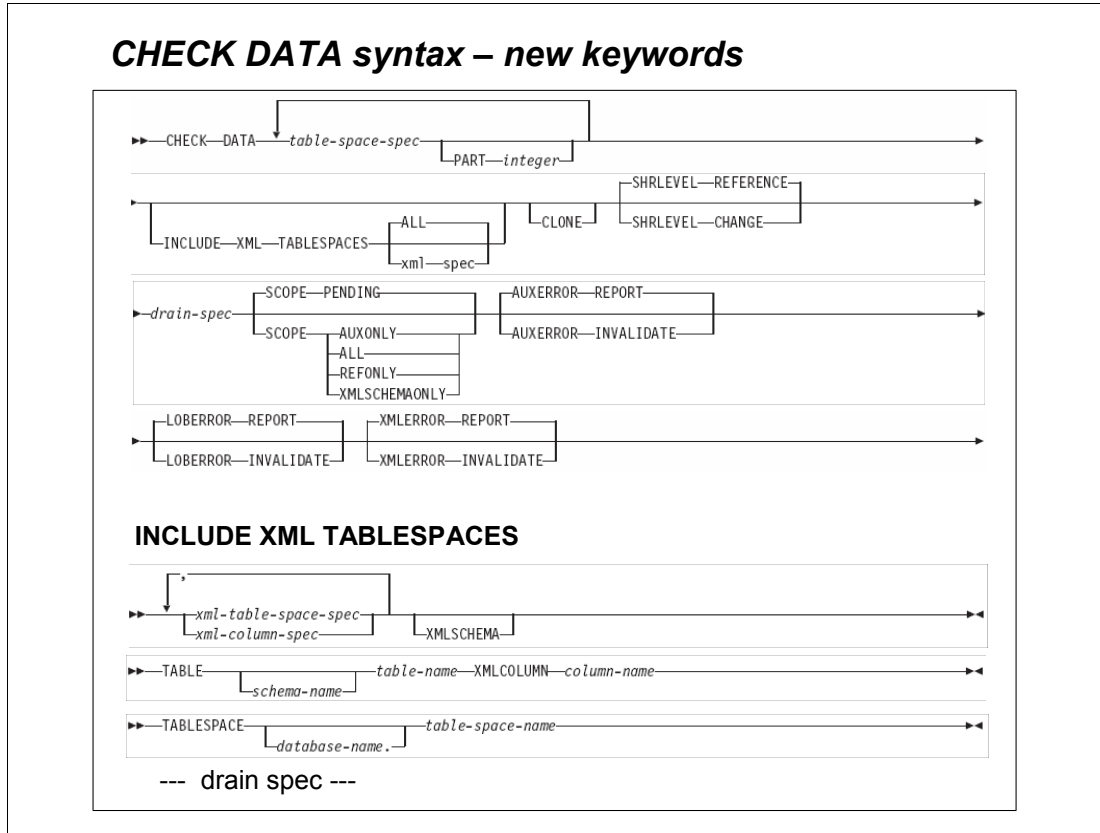


Figure 9-1 CHECK DATA syntax - new keywords

Options INCLUDE XML TABLESPACES and SCOPE XMLSCHEMAONLY are new.

If you specify the INCLUDE XML TABLESPACES option, CHECK DATA can check the structural integrity of XML documents. CHECK DATA can verify the following items for XML objects:

- ▶ All rows in an XML column exist in the XML table space.
- ▶ All documents in the XML table space are structurally valid.
- ▶ Each index entry in the NODEID index has a corresponding XML document.
- ▶ Each XML document in the XML table space has corresponding entries in the NODEID index.
- ▶ Each entry in the DOCID column in the base table space has a corresponding entry in the NODEID index over the XML table space, if the XML column is not null.
- ▶ Each entry in the NODEID index contains a corresponding value in the DOCID column.
- ▶ If an XML column has an XML type modifier, all XML documents in the column are valid with respect to at least one XML schema that is associated with the XML type modifier.

If the base table space is not consistent with any related XML table spaces, or a problem is found during any of the previously listed checks, CHECK DATA reports the error.

For XML checking, the default behavior of CHECK DATA is to check only the consistency between each XML column and its NODEID index. However, you can modify the scope of checking by specifying combinations of the CHECK DATA SCOPE keyword and the INCLUDE

XML TABLESPACES keyword. For LOBs, CHECK DATA utility checks for the consistency between the base table and the auxiliary index.

Table 9-1 is a reference table for the CHECK DATA invocation based on the combination of the various options which are applicable for LOBs as well.

Table 9-1 CHECK DATA invocation

CHECK DATA base table space(s)	INCLUDE XML TABLESPACES	XMLSCHEMA	SCOPE	Structure check	Schema validation	XML checks	LOB checks
X			ALL/AUXONLY	-	-	Yes	Yes
X			REFONLY	-	-	-	-
X			XMLSCHEMA ONLY	-	Yes Default: INCLUDE XML TABLESPACES ALL	-	-
X			PENDING	-	-	Yes Base table spaces in CHKP/ACHKP	Yes Base table spaces in CHKP/ACHKP
X	X		ALL/AUXONLY	Yes Specified XML table spaces only	-	Yes All XML table spaces	Yes All LOB table spaces
X	X		REFONLY	-	-	-	-
X	X		XMLSCHEMA ONLY	-	Yes Specified XML table spaces only	-	-
X	X		PENDING	-	Yes Specified XML table spaces in CHKP	Yes Base table spaces in CHKP/ACHKP	Yes Base table spaces in CHKP/ACHKP
X	X	X	ALL/AUXONLY	Yes Specified XML table spaces only	Yes Specified XML table spaces only	Yes All XML table spaces	Yes All LOB table spaces
X	X	X	REFONLY	-	-	-	-
X	X	X	XMLSCHEMA ONLY	-	Yes Specified XML table spaces only	-	-
X	X	X	PENDING	-	Yes Specified XML table spaces in CHKP	Yes Base table spaces in CHKP/ACHKP	Yes Base table spaces in CHKP/ACHKP

The 'XML checks' column indicates CHECK DATA utility checks only for the consistency between the base table and the NODEID index.

The 'LOB checks' column indicates CHECK DATA utility checks for the consistency between the base table and the auxiliary index.

Example 9-1 shows an example of CHECK DATA utility control statement for our scenario.

Example 9-1 CHECK DATA example

```
CHECK DATA TABLESPACE DSN00242.BKRTORCS SCOPE ALL
```



```
INCLUDE XML TABLESPACES(TABLE BK_TO_CSTMR_STMT XMLCOLUMN BK_TO_CSTMR_STMT) XMLSCHEMA
```

Table space DSN00242.BKRTORCS contains table BK_TO_CSTMR_STMT with XML column BK_TO_CSTMR_STMT which has an XML type modifier. If you specify the statement as shown in Example 9-1, CHECK DATA checks LOB relationships, the base table space, XML relationships, and the structural integrity of XML documents for column BK_TO_CSTMR_STMT, and does XML schema validation on the documents for column BK_TO_CSTMR_STMT.

You can use the option SHRLEVEL REFERENCE or SHRELEVEL CHANGE.

Figure 9-2 shows considerations with option SHRLEVEL REFERENCE.

- ... **SHRLEVEL REFERENCE and AUXERROR/XMLERROR REPORT**
 - If no problem found, remove CHKP or ACHKP from table spaces
 - If problem found:
 - DOCID of XML document is printed in the job output
 - No further action
- ... **SHRLEVEL REFERENCE and AUXERROR/XMLERROR INVALIDATE**
 - If no problem found, remove CHKP or ACHKP from table spaces
 - If problem found:
 - DOCID of XML document is printed in the job output
 - Exception tables automatically generated for XML column (schema validation only)
 - Affected XML documents moved to XML exception table (schema validation only)
 - Corrupted XML document deleted from XML table space
 - Index entries for corrupted XML documents removed from NODEID index
 - Invalid bit set in the XML indicator in the base table space
 - Value index(es) not touched/checked by CHECK DATA

Figure 9-2 CHECK DATA - SHRLEVEL REFERENCE considerations

Figure 9-3 shows considerations with option SHRLEVEL CHANGE.

- ... **SHRLEVEL CHANGE in general**
 - Utility creates shadow copies of all table and index spaces
 - Shadow copies discarded at the end of utility execution
- ... **SHRLEVEL CHANGE and AUXERROR/XMLERROR REPORT**
 - If no problem found, CHKP or ACHKP remain on table spaces
 - If problem found:
 - DOCID of XML document is printed in the job output
 - No further action

Figure 9-3 CHECK DATA - SHRLEVEL CHANGE considerations (1 of 2)

Example 9-4 shows more considerations with SHRLEVEL CHANGE.

- If no problem found, CHKP or ACHKP remain on table spaces
- If problem found:
 - DOCID of XML document is printed in the job output
 - For each corrupted XML document CHECK DATA creates REPAIR LOCATE... DOCID.. DELETE
 - Message issued telling you to run REBUILD INDEX on NODEID index if a problem is found
 - **NO RBDP set on NODEID index**
 - REPAIR LOCATE ... RID... REPLACE statements generated to invalidate entry.
 - Value indexes not touched/checked by CHECK DATA
- CHECK DATA generates REPAIR statements for XML documents which are not valid according to the defined XML type modifier also:
 - REPAIR LOCATE ... DOCID ... DELETE
 - REPAIR LOCATE ... RID ... REPLACE
- Run REPAIR
 - Delete corrupted XML documents from XML table space
`REPAIR LOCATE TABLESPACE "DSN00155 ". "XBKR0000"
DOCID 1 DELETE SHRLEVEL CHANGE`
 - Set invalid bit in the XML indicator column in the base table space
`REPAIR LOCATE TABLESPACE "DSN00155 ". "BKR TORCS "
RID X'0123456789ABCDEF'
VERIFY OFFSET 28 DATA X'ABCD'
REPLACE OFFSET 28 DATA X'1234'`

Figure 9-4 CHECK DATA - SHRLEVEL CHANGE considerations (2 of 2)

Refer to 10.5, “Diagnostics” on page 238 for examples of invoking the CHECK DATA utility when diagnosing problems with XML data.

9.2 CHECK INDEX

You can use the CHECK INDEX utility to check XML indexes, document ID indexes, and NODEID indexes. You do not need to specify any additional keywords. You cannot specify the DOCID and NODEID indexes in a single CHECK INDEX control statement because they belong to two different table spaces. Specify the control statement using two CHECK INDEX statements in the utility run as shown in Example 9-2.

Example 9-2 CHECK INDEX utility JCL and output

```
//XMLR4LD JOB (999,POK), 'DB0B', CLASS=A,
// MSGCLASS=T, NOTIFY=XMLR4, TIME=1440, REGION=0M
/*JOBPARM SYSAFF=SC63, L=9999
// JCLLIB ORDER=(DB0BM.PROCLIB)
//DSNUPROC.SORTWK01 DD DSN=XMLR4.SORTWK01,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK02 DD DSN=XMLR4.SORTWK02,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK03 DD DSN=XMLR4.SORTWK03,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
```

```

//      UNIT=SYSDA
//DSNUPROC.SORTWK04 DD DSN=XMLR4.SORTWK04,
//      DISP=(MOD,DELETE,CATLG),
//      SPACE=(16384,(20,20),,,ROUND),
//      UNIT=SYSDA
//DSNUPROC.SYSUT1 DD DSN=XMLR4.SYSUT1,
//      DISP=(MOD,DELETE,CATLG),
//      SPACE=(16384,(20,20),,,ROUND),
//      UNIT=SYSDA
//DSNUPROC.SYSIN DD *
CHECK INDEX (XMLR4.I_DOCIDBK_TO_CSTMR_STMT)
CHECK INDEX (XMLR4.I_NODEIDXBK_TO_CSTMR_STMT)

1DSNU000I    314 15:32:32.95 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
  DSNU1044I  314 15:32:32.97 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I    314 15:32:32.98 DSNUGUTC - CHECK INDEX(XMLR4.I_DOCIDBK_TO_CSTMR_STMT)
  DSNU395I   314 15:32:32.99 DSNUKPIK - INDEXES WILL BE CHECKED IN PARALLEL, NUMBER OF
TASKS = 3
  DSNU701I   -DBOB 314 15:32:32.99 DSNUKIUL - 1 INDEX ENTRIES UNLOADED FROM
'DSN00242.BKRTORCS'
  DSNU705I   314 15:32:33.00 DSNUKPIK - UNLOAD PHASE COMPLETE - ELAPSED TIME=00:00:00
  DSNU719I   314 15:32:33.07 DSNUKPIK - 1 ENTRIES CHECKED FOR INDEX
'XMLR4.I_DOCIDBK_TO_CSTMR_STMT'
  DSNU720I   314 15:32:33.07 DSNUKPIK - SORTCHK PHASE COMPLETE, ELAPSED TIME=00:00:00
  DSNU719I   -DBOB 314 15:32:33.07 DSNUKTER - 1 ENTRIES CHECKED FOR INDEX
'XMLR4.I_DOCIDBK_TO_CSTMR_STMT'
  DSNU380I   -DBOB 314 15:32:33.07 DSNUGRSX - TABLESPACE DSN00242.BKRTORCS PARTITION 1 IS IN
COPY PENDING

ODSNU050I    314 15:32:33.07 DSNUGUTC - CHECK INDEX(XMLR4.I_NODEIDXBK_TO_CSTMR_STMT)
  DSNU395I   314 15:32:33.08 DSNUKPIK - INDEXES WILL BE CHECKED IN PARALLEL, NUMBER OF
TASKS = 3
  DSNU701I   -DBOB 314 15:32:33.08 DSNUKIUL - 1 INDEX ENTRIES UNLOADED FROM
'XMLR4.I_NODEIDXBK_TO_CSTMR_STMT'
  DSNU705I   314 15:32:33.09 DSNUKPIK - UNLOAD PHASE COMPLETE - ELAPSED TIME=00:00:00
  DSNU719I   314 15:32:33.12 DSNUKPIK - 1 ENTRIES CHECKED FOR INDEX
'XMLR4.I_NODEIDXBK_TO_CSTMR_STMT'
  DSNU720I   314 15:32:33.12 DSNUKPIK - SORTCHK PHASE COMPLETE, ELAPSED TIME=00:00:00
  DSNU719I   -DBOB 314 15:32:33.12 DSNUKTER - 1 ENTRIES CHECKED FOR INDEX
'XMLR4.I_NODEIDXBK_TO_CSTMR_STMT'
  DSNU380I   -DBOB 314 15:32:33.12 DSNUGRSX - TABLESPACE DSN00242.XBKRO000 PARTITION 1 IS IN
COPY PENDING

```

NOTE: You can specify the CHECK INDEX control statement as shown below:

```
CHECK INDEX(ALL) TABLESPACE DSN00242.BKRTORCS
```

```
CHECK INDEX(ALL) TABLESPACE DSN00242.XBKRO000
```

The advantage with this approach is any user defined indexes are also checked.

After you run the CHECK INDEX utility, you might need to correct XML data.

To correct XML data, based on the CHECK INDEX output, perform one of the actions shown in Table 9-2.

Table 9-2 Action to be taken based on CHECK INDEX output

Problem	Action
Problem with a document ID index	<ol style="list-style-type: none"> 1. Confirm that the base table space is at the correct level. 2. Rebuild the index.
Problem with an XML table space for a NODEID index or an XML index and the index is correct	Run REPAIR LOCATE RID DELETE to remove the orphan row.
Problem with an XML table space for a NODEID index or an XML index and the index is incorrect	Run REBUILD INDEX or RECOVER INDEX to rebuild the index.
Problem with an XML index over an XML table space	Run REBUILD INDEX to rebuild the index. Restriction: Do not run REPAIR LOCATE RID DELETE to remove orphan rows unless the NODEID index does not represent the same row and the base table space does not use the document ID index.

Note: CHECK INDEX of an XML index cannot run if REBUILD INDEX, REORG INDEX, or RECOVER is being run on that index because CHECK INDEX needs access to the NODEID index. CHECK INDEX SHRLEVEL CHANGE cannot run two jobs concurrently for two different indexes that are in the same table space.

9.3 COPY

You can use the COPY utility to copy XML objects. You do not need to specify any additional keywords. When you specify that DB2 is to copy a table space with XML columns, DB2 does not automatically copy any related XML table spaces or indexes. You must explicitly specify the XML objects that you want to copy.

The COPY utility control statement in Example 9-3 specifies that DB2 is to copy base table space DSN00242.BKRTORCS and the XML table space DSN00242.XBKR0000.

Example 9-3 COPY utility JCL for taking full image copy and output

```
//XMLR4LD JOB (999,POK),'DB0B',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=OM
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DB0BM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DB0B,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
TEMPLATE A DSN(&DB..&SN..&IC..D&DATE..T&TIME..COPY)
COPY TABLESPACE DSN00242.BKRTORCS COPYDDN(A)
      TABLESPACE DSN00242.XBKR0000 COPYDDN(A)

1DSNU000I    314 18:12:48.11 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I    314 18:12:48.14 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I    314 18:12:48.14 DSNUGUTC - TEMPLATE A
DSN(&DB..&SN..&IC..D&DATE..T&TIME..COPY)
DSNU1035I    314 18:12:48.14 DSNUJTDR - TEMPLATE STATEMENT PROCESSED SUCCESSFULLY
```

```

ODSNU050I  314 18:12:48.14 DSNUGUTC - COPY TABLESPACE DSN00242.BKRTORCS COPYDDN(A)
TABLESPACE DSN00242.XBKRO000 COPYDDN(A)
DSNU1038I  314 18:12:48.19 DSNUGDYN - DATASET ALLOCATED.  TEMPLATE=A
          DDNAME=SYS00001
          DSN=DSN00242.BKRTORCS.F.D2010314.T231248.COPY
DSNU400I   314 18:12:48.24 DSNUBBID - COPY PROCESSED FOR TABLESPACE DSN00242.BKRTORCS
          NUMBER OF PAGES=3
          AVERAGE PERCENT FREE SPACE PER PAGE = 32.66
          PERCENT OF CHANGED PAGES = 0.00
          ELAPSED TIME=00:00:00
DSNU1038I  314 18:12:48.27 DSNUGDYN - DATASET ALLOCATED.  TEMPLATE=A
          DDNAME=SYS00002
          DSN=DSN00242.XBKRO000.F.D2010314.T231248.COPY
DSNU400I   314 18:12:48.30 DSNUBBID - COPY PROCESSED FOR TABLESPACE DSN00242.XBKRO000
          NUMBER OF PAGES=3
          AVERAGE PERCENT FREE SPACE PER PAGE = 19.66
          PERCENT OF CHANGED PAGES = 0.00
          ELAPSED TIME=00:00:00
DSNU428I  -DB0B 314 18:12:48.31 DSNUBAFI - DB2 IMAGE COPY SUCCESSFUL FOR TABLESPACE
DSN00242.BKRTORCS
DSNU428I  -DB0B 314 18:12:48.31 DSNUBAFI - DB2 IMAGE COPY SUCCESSFUL FOR TABLESPACE
DSN00242.XBKRO000

```

Both full and incremental image copies are supported for an XML table space, as well as the SHRLEVEL REFERENCE, SHRLEVEL CHANGE, CONCURRENT, and FLASHCOPY options.

To demonstrate using COPY utility to take an incremental image copy, the XML document is modified as shown in Figure 9-5.

```

SELECT XMLSERIALIZE(
  XMLQUERY(
    'declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    /Document/BkToCstmrStmt/Stmt/Bal/Amt[../Tp/CdOrPrtry/Cd="CLBD"] '
    PASSING BK_TO_CSTMR_STMT) AS CLOB(500))
FROM BK_TO_CSTMR_STMT;

<Amt xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02"
Ccy="SEK">435678.50</Amt>
DSNE610I NUMBER OF ROWS DISPLAYED IS 1

UPDATE BK_TO_CSTMR_STMT
SET BK_TO_CSTMR_STMT = XMLMODIFY (
  'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
  replace value of node
  /Document/BkToCstmrStmt/Stmt/Bal/Amt[../Tp/CdOrPrtry/Cd="CLBD"]
  with "900000"')
WHERE MSG_ID IS NULL ;

DSNE615I NUMBER OF ROWS AFFECTED IS 1

SELECT XMLSERIALIZE(
  XMLQUERY(
    'declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    /Document/BkToCstmrStmt/Stmt/Bal/Amt[../Tp/CdOrPrtry/Cd="CLBD"] '
    PASSING BK_TO_CSTMR_STMT) AS CLOB(500))
FROM BK_TO_CSTMR_STMT;

<Amt xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02" Ccy="SEK">900000</Amt>
DSNE610I NUMBER OF ROWS DISPLAYED IS 1

```

Figure 9-5 Make a partial update to the XML document

The COPY utility control statement in Example 9-4 specifies that DB2 is to take an incremental image copy of base table space DSN00242.BKRTORCS and the XML table space DSN00242.XBKR0000.

Example 9-4 COPY utility JCL for taking incremental image copy and output

```

//XMLR4LD JOB (999,POK),'DB0B',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DB0BM.PROCLIB)
//DSNUPROC.SYSIN DD *
TEMPLATE A DSN(&DB..&SN..&IC..&DATE..T&TIME..COPY)
COPY TABLESPACE DSN00242.BKRTORCS COPYDDN(A) FULL NO
      TABLESPACE DSN00242.XBKR0000 COPYDDN(A) FULL NO

1DSNU000I 314 18:14:25.22 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 314 18:14:25.25 DSNUGTIS - PROCESSING SYSIN AS EBCDIC

```

```

ODSNU050I 314 18:14:25.26 DSNUGUTC - TEMPLATE A
DSN(&DB..&SN..&IC..D&DATE..T&TIME..COPY)
DSNU1035I 314 18:14:25.26 DSNUJTDR - TEMPLATE STATEMENT PROCESSED SUCCESSFULLY

ODSNU050I 314 18:14:25.26 DSNUGUTC - COPY TABLESPACE DSN00242.BKRTORCS COPYDDN(A) FULL
NO TABLESPACE DSN00242.XBKRO000 COPYDDN(A) FULL NO
DSNU1038I 314 18:14:25.30 DSNUGDYN - DATASET ALLOCATED. TEMPLATE=A
DDNAME=SYS00001
DSN=DSN00242.BKRTORCS.I.D2010314.T231425.COPY
DSNU400I 314 18:14:25.34 DSNUBBID - COPY PROCESSED FOR TABLESPACE DSN00242.BKRTORCS
NUMBER OF PAGES=3
AVERAGE PERCENT FREE SPACE PER PAGE = 32.66
PERCENT OF CHANGED PAGES = 5.88
ELAPSED TIME=00:00:00
DSNU1038I 314 18:14:25.37 DSNUGDYN - DATASET ALLOCATED. TEMPLATE=A
DDNAME=SYS00002
DSN=DSN00242.XBKRO000.I.D2010314.T231425.COPY
DSNU400I 314 18:14:25.41 DSNUBBID - COPY PROCESSED FOR TABLESPACE DSN00242.XBKRO000
NUMBER OF PAGES=3
AVERAGE PERCENT FREE SPACE PER PAGE = 6.33
PERCENT OF CHANGED PAGES = 4.44
ELAPSED TIME=00:00:00
DSNU428I -DB0B 314 18:14:25.42 DSNUBAFI - DB2 IMAGE COPY SUCCESSFUL FOR TABLESPACE
DSN00242.BKRTORCS
DSNU428I -DB0B 314 18:14:25.42 DSNUBAFI - DB2 IMAGE COPY SUCCESSFUL FOR TABLESPACE
DSN00242.XBKRO000

```

Unless either the CONCURRENT option or the FLASHCOPY option is specified, COPY does not copy empty or unformatted data pages of an XML table space.

To copy an XML table space with a base table space that has the NOT LOGGED attribute, all associated XML table spaces must also have the NOT LOGGED attribute. The XML table space acquires this NOT LOGGED attribute by being linked to the logging attribute of its associated base table space. You cannot independently alter the logging attribute of an XML table space.

If the LOG column of the SYSIBM.SYSTABLESPACE record for an XML table space has the value of "X", the logging attributes of the XML table space and its base table space are linked, and that the logging attribute of both table spaces is NOT LOGGED. To break the link, alter the logging attribute of the base table space back to LOGGED, and the logging attribute of both table spaces are changed back to LOGGED

9.4 COPYTOCOPY

You can use the COPYTOCOPY utility to copy existing copies of the XML objects. You do not need to specify any additional keywords.

The COPYTOCOPY utility control statement in Example 9-5 specifies that DB2 is to make primary and backup copies for the remote site for the XML table space DSN00242.XBKRO000 using the last full image copy that was created.

Example 9-5 COPYTOCOPY utility JCL and output

```

//XMLR4LD JOB (999,POK),'DB0B',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DB0BM.PROCLIB)

```

```
//UTIL EXEC DSNUPROC,SYSTEM=DB0B,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
TEMPLATE A DSN(&DB.&SN.&IC.&PB.&DATE.&TIME.&COPY)
COPYTOCOPY TABLESPACE DSN00242.XBKRO000
FROMLASTFULLCOPY
RECOVERYDDN(A,A)

1DSNU000I 314 19:22:30.00 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
  DSNU1044I 314 19:22:30.04 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 314 19:22:30.04 DSNUGUTC - TEMPLATE A
DSN(&DB.&SN.&IC.&PB.&DATE.&TIME.&COPY)
  DSNU1035I 314 19:22:30.04 DSNUJTDR - TEMPLATE STATEMENT PROCESSED SUCCESSFULLY
1DSNU000I 314 19:22:31.81 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
  DSNU1044I 314 19:22:31.83 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 314 19:22:31.83 DSNUGUTC - COPYTOCOPY TABLESPACE DSN00242.XBKRO000
FROMLASTFULLCOPY RECOVERYDDN(A,A)
DSNU1403I 314 19:22:31.87 DSNU2BCC - LOCAL SITE PRIMARY DATA SET
DSN00242.XBKRO000.F.D2010314.T233001.COPY WITH
      START_RBA 0000696BCB4C IS IN USE BY COPYTOCOPY
      FOR TABLESPACE DSN00242.XBKRO000
DSNU1404I 314 19:22:31.88 DSNU2BDR - COPYTOCOPY PROCESSING COMPLETED FOR
      TABLESPACE DSN00242.XBKRO000
      ELAPSED TIME = 00:00:00
      NUMBER OF PAGES COPIED=3
DSNU1406I 314 19:22:31.89 DSNU2BDR - COPYTOCOPY COMPLETED. ELAPSED TIME = 00:00:00
```

9.5 EXEC SQL

The EXEC SQL utility control statement declares cursors or executes dynamic SQL statements. You can use this utility as part of the DB2 cross-loader function of the LOAD utility.

The cross-loader function enables you to use a single LOAD job to transfer data from one location to another location or from one table to another table at the same location. You can use either a local server or any DRDA-compliant remote server as a data input source for populating your tables. Your input can even come from other sources besides DB2 for z/OS; you can use IBM Information Integrator Federation feature for access to data from sources as diverse as Oracle and Sybase, as well as the entire DB2 family of database servers.

Note: You cannot declare a cursor that includes XML data. Thus, you cannot use the DB2 family cross-loader function to transfer data from XML columns. However, you can declare a cursor on a table with XML columns if the cursor does not include any XML columns.

For example, suppose that you create the following table with an XML column:

```
CREATE TABLE BK_TO_CSTMR_STMT
(MSG_ID VARCHAR(35),
MSG_CRE_DT_TM TIMESTAMP WITH TIMEZONE,
BK_TO_CSTMR_STMT XML NOT NULL)
```

You cannot declare the following cursor, because it includes XML data in the BK_TO_CSTMR_STMT column:

```
EXEC SQL
DECLARE C1 CURSOR FOR SELECT * FROM BK_TO_CSTMR_STMT
```



```
END-EXEC
```

However, you can declare a cursor that includes non-XML columns, as in the following example:

```
EXEC SQL
DECLARE C2 CURSOR FOR SELECT MSG_ID FROM BK_TO_CSTMRT
END-EXEC
```

9.6 LISTDEF

When you create object lists with the LISTDEF utility, you can specify whether you want related XML objects to be included or excluded.

Use the following keywords to indicate the objects to include or exclude:

- ▶ **ALL** Base and XML objects (This keyword is the default.)
- ▶ **BASE** Base objects only
- ▶ **XML** XML objects only

For example, the LISTDEF statements in Table 9-3 generate the indicated lists.

Table 9-3 Example LISTDEF statements

LISTDEF statement followed by objects that are included in the list
<pre>LISTDEF LISTALL INCLUDE TABLESPACES DATABASE DSN00242 INCLUDE INDEXSPACES DATABASE DSN00242</pre> <ul style="list-style-type: none"> ▶ All tables spaces in the DSN00242 database, including XML table spaces ▶ All index spaces in the DSN00242 database
<pre>LISTDEF LISTXML INCLUDE TABLESPACES DATABASE DSN00242 XML INCLUDE INDEXSPACES DATABASE DSN00242 XML</pre> <ul style="list-style-type: none"> ▶ All XML table spaces in the DSN00242 database ▶ All XML index spaces in the DSN00242 database
<pre>LISTDEF LIST INCLUDE TABLESPACES DATABASE DSN00242 ALL INCLUDE INDEXSPACES DATABASE DSN00242 ALL EXCLUDE INDEXSPACES DATABASE DSN00242 XML</pre> <ul style="list-style-type: none"> ▶ All tables spaces in the DSN00242 database, including XML table spaces ▶ All index spaces in the DSN00242 database except for XML index spaces

Example 9-6 shows the JCL for the LISTDEF utility for the first LISTDEF statement in Table 9-3 and the output of the utility run.

Example 9-6 JCL for LISTDEF utility and output (1 of 3)

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=' '
//DSNUPROC.SYSIN DD *
OPTIONS PREVIEW
LISTDEF LISTALL INCLUDE TABLESPACES DATABASE DSN00242
INCLUDE INDEXSPACES DATABASE DSN00242
```

```

1DSNU000I 314 20:59:55.38 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
  DSNU1044I 314 20:59:55.41 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 314 20:59:55.41 DSNUGUTC - OPTIONS PREVIEW
  DSNU1000I 314 20:59:55.41 DSNUZODR - PROCESSING CONTROL STATEMENTS IN PREVIEW MODE
  DSNU1035I 314 20:59:55.41 DSNUZODR - OPTIONS STATEMENT PROCESSED SUCCESSFULLY
ODSNU050I 314 20:59:55.41 DSNUGUTC - LISTDEF LISTALL INCLUDE TABLESPACES DATABASE
DSN00242 INCLUDE INDEXSPACES DATABASE DSN00242
  DSNU1035I 314 20:59:55.41 DSNUILDR - LISTDEF STATEMENT PROCESSED SUCCESSFULLY
  DSNU1020I -DBOB 314 20:59:55.41 DSNUILSA - EXPANDING LISTDEF LISTALL
  DSNU1021I -DBOB 314 20:59:55.41 DSNUILSA - PROCESSING INCLUDE CLAUSE DATABASE DSN00242.
  DSNU1022I -DBOB 314 20:59:55.42 DSNUILSA - CLAUSE IDENTIFIES 2 OBJECTS
  DSNU1021I -DBOB 314 20:59:55.42 DSNUILSA - PROCESSING INCLUDE CLAUSE DATABASE DSN00242.
  DSNU1022I -DBOB 314 20:59:55.43 DSNUILSA - CLAUSE IDENTIFIES 2 OBJECTS
  DSNU1023I -DBOB 314 20:59:55.43 DSNUILSA - LISTDEF LISTALL CONTAINS 4 OBJECTS
  DSNU1010I 314 20:59:55.43 DSNUGPVV - LISTDEF LISTALL EXPANDS TO THE FOLLOWING OBJECTS:
    LISTDEF LISTALL -- 00000004 OBJECTS
    INCLUDE TABLESPACE DSN00242.BKRTRORCS
    INCLUDE TABLESPACE DSN00242.XBKRO000
    INCLUDE INDEXSPACE DSN00242.IRDOCIDB
    INCLUDE INDEXSPACE DSN00242.IRNODEID

```

Example 9-7 shows the JCL for the LISTDEF utility for the second LISTDEF statement in Table 9-3 on page 193 and the output of the utility run.

Example 9-7 JCL for LISTDEF utility and output (2 of 3)

```

//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
OPTIONS PREVIEW
LISTDEF LISTXML INCLUDE TABLESPACES DATABASE DSN00242 XML
          INCLUDE INDEXSPACES DATABASE DSN00242 XML

1DSNU000I 314 21:00:53.19 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
  DSNU1044I 314 21:00:53.22 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 314 21:00:53.22 DSNUGUTC - OPTIONS PREVIEW
  DSNU1000I 314 21:00:53.22 DSNUZODR - PROCESSING CONTROL STATEMENTS IN PREVIEW MODE
  DSNU1035I 314 21:00:53.22 DSNUZODR - OPTIONS STATEMENT PROCESSED SUCCESSFULLY
ODSNU050I 314 21:00:53.22 DSNUGUTC - LISTDEF LISTXML INCLUDE TABLESPACES DATABASE
DSN00242 XML INCLUDE INDEXSPACES DATABASE DSN00242 XML
  DSNU1035I 314 21:00:53.22 DSNUILDR - LISTDEF STATEMENT PROCESSED SUCCESSFULLY
  DSNU1020I -DBOB 314 21:00:53.22 DSNUILSA - EXPANDING LISTDEF LISTXML
  DSNU1021I -DBOB 314 21:00:53.22 DSNUILSA - PROCESSING INCLUDE CLAUSE DATABASE DSN00242.
  DSNU1022I -DBOB 314 21:00:53.22 DSNUILSA - CLAUSE IDENTIFIES 1 OBJECTS
  DSNU1021I -DBOB 314 21:00:53.22 DSNUILSA - PROCESSING INCLUDE CLAUSE DATABASE DSN00242.
  DSNU1022I -DBOB 314 21:00:53.22 DSNUILSA - CLAUSE IDENTIFIES 1 OBJECTS
  DSNU1023I -DBOB 314 21:00:53.22 DSNUILSA - LISTDEF LISTXML CONTAINS 2 OBJECTS
  DSNU1010I 314 21:00:53.22 DSNUGPVV - LISTDEF LISTXML EXPANDS TO THE FOLLOWING OBJECTS:
    LISTDEF LISTXML -- 00000002 OBJECTS
    INCLUDE TABLESPACE DSN00242.XBKRO000
    INCLUDE INDEXSPACE DSN00242.IRNODEID

```

Example 9-8 shows the JCL for the LISTDEF utility for the third LISTDEF statement in Table 9-3 on page 193 and the output of the utility run.

Example 9-8 JCL for LISTDEF utility and output (3 of 3)

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
OPTIONS PREVIEW
LISTDEF LIST INCLUDE TABLESPACES DATABASE DSN00242 ALL
          INCLUDE INDEXSPACES DATABASE DSN00242 ALL
          EXCLUDE INDEXSPACES DATABASE DSN00242 XML
1DSNU000I 314 21:02:09.82 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
  DSNU1044I 314 21:02:09.85 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
0DSNU050I 314 21:02:09.85 DSNUGUTC - OPTIONS PREVIEW
  DSNU1000I 314 21:02:09.85 DSNUZODR - PROCESSING CONTROL STATEMENTS IN PREVIEW MODE
  DSNU1035I 314 21:02:09.85 DSNUZODR - OPTIONS STATEMENT PROCESSED SUCCESSFULLY
0DSNU050I 314 21:02:09.85 DSNUGUTC - LISTDEF LIST INCLUDE TABLESPACES DATABASE DSN00242
ALL INCLUDE INDEXSPACES
  DATABASE DSN00242 ALL EXCLUDE INDEXSPACES DATABASE DSN00242 XML
  DSNU1035I 314 21:02:09.85 DSNUILDR - LISTDEF STATEMENT PROCESSED SUCCESSFULLY
  DSNU1020I -DBOB 314 21:02:09.85 DSNUILSA - EXPANDING LISTDEF LIST
  DSNU1021I -DBOB 314 21:02:09.85 DSNUILSA - PROCESSING INCLUDE CLAUSE DATABASE DSN00242.
  DSNU1022I -DBOB 314 21:02:09.86 DSNUILSA - CLAUSE IDENTIFIES 2 OBJECTS
  DSNU1021I -DBOB 314 21:02:09.86 DSNUILSA - PROCESSING INCLUDE CLAUSE DATABASE DSN00242.
  DSNU1022I -DBOB 314 21:02:09.86 DSNUILSA - CLAUSE IDENTIFIES 2 OBJECTS
  DSNU1021I -DBOB 314 21:02:09.86 DSNUILSA - PROCESSING EXCLUDE CLAUSE DATABASE DSN00242.
  DSNU1022I -DBOB 314 21:02:09.86 DSNUILSA - CLAUSE IDENTIFIES 1 OBJECTS
  DSNU1023I -DBOB 314 21:02:09.86 DSNUILSA - LISTDEF LIST CONTAINS 3 OBJECTS
  DSNU1010I 314 21:02:09.86 DSNUGPVV - LISTDEF LIST EXPANDS TO THE FOLLOWING OBJECTS:
    LISTDEF LIST -- 00000003 OBJECTS
      INCLUDE TABLESPACE DSN00242.BKRTORCS
      INCLUDE TABLESPACE DSN00242.XBKRO000
      INCLUDE INDEXSPACE DSN00242.IRDOCIDB
```

9.7 LOAD

You can load data containing XML columns with one of two methods.

- ▶ The XML column can be loaded from the input record. XML column value can be placed in the INPUT record with or without any other loading column values. The input record can be in delimited or non-delimited format.

For a non-delimited format, the XML column is treated like a variable character with a 2-byte length preceding the XML value. For a delimited format there are no length bytes present. If the input record is in spanned record format, specify the FORMAT SPANNED YES option.

- ▶ The XML column can be loaded from a separate file whether the XML column length is less than 32 KB or not.

To load data into a base table that has XML columns:

1. Create input data sets to ensure that you use the appropriate format:
 - If the data set is in delimited format, ensure that the XML input fields follow the standard LOAD utility delimited format.

- If the data set is not in delimited format, specify the XML input fields similar to the way that you specify VARCHAR input. Specify the length of the field in a 2-byte binary field that precedes the data.
2. Create a LOAD utility control statement.
- To load XML directly from input record, specify XML as the input field type. XML is the only acceptable field type and data type conversion is not supported. Do not specify DEFAULTIF.
- If you want the whitespace to be preserved in the XML data, also specify the keywords PRESERVE WHITESPACE. By default, LOAD strips the whitespace.
- When data in the binary XML format is loaded into a table, and PRESERVE WHITESPACE is not specified, DB2 strips whitespace only when the input data contains whitespace tags.
- To load XML from a file, specify CHAR or VARCHAR along with either BLOBF, CLOBF or DBCLOBF to indicate that the input column contains a filename from which a BLOBF, CLOBF or DBCLOBF is to be loaded to the XML column.
3. Submit the utility control statement.

Note: When you load XML documents into a table, and the XML value cannot be cast to the type that you specified when you created the index, the value is ignored without any warnings or errors, and the document is inserted into the table.

When you insert XML documents into a table with XML indexes that are of type DECFLOAT, the values might be rounded when they are inserted. If the index is unique, the rounding might cause duplicates even if the original values are not exactly the same.

Example 9-9 shows the JCL for the LOAD utility and the output of the utility run.

Example 9-9 LOAD utility JCL and output

```
//XMLR4LD JOB (999,POK), 'DBOB', CLASS=A,
// MSGCLASS=T, NOTIFY=&SYSUID, TIME=1440, REGION=0M
/*JOBPARM SYSAFF=SC63, L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//LOAD1 EXEC DSNUPROC, SYSTEM=DBOB, UID=' '
//SYSREC DD DSN=XMLR4.BKSTMT.XMLDATA, DISP=SHR
//SYSERR DD DSN=XMLR4.LOAD.SYSERR,
// DISP=(MOD,DELETE,CATLG), UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//SYSDISC DD DSN=XMLR4.LOAD.SYSDISC,
// DISP=(MOD,DELETE,CATLG), UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//SYSMAP DD DSN=XMLR4.LOAD.SYSMAP,
// DISP=(MOD,DELETE,CATLG), UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//SYSUT1 DD DSN=XMLR4.LOAD.SYSUT1,
// DISP=(MOD,DELETE,CATLG), UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//UTPRINT DD SYSOUT=*
//SORTOUT DD DSN=XMLR4.LOAD.SORTOUT,
// DISP=(MOD,DELETE,CATLG), UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//DSNUPROC.SYSIN DD *
LOAD DATA REPLACE
      INTO TABLE XMLR4.BK_TO_CSTMR_STMT
      (BK_TO_CSTMR_STMT POSITION(1) XML PRESERVE WHITESPACE)
```

```

1DSNU000I    300 20:39:12.45 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = XMLR4.XMLR4LD
DSNU1044I    300 20:39:12.47 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I    300 20:39:12.48 DSNUGUTC - LOAD DATA REPLACE
DSNU650I    -DBOB 300 20:39:12.48 DSNURWI - INTO TABLE XMLR4.BK_TO_CSTMR_STMT
DSNU650I    -DBOB 300 20:39:12.48 DSNURWI - (BK_TO_CSTMR_STMT POSITION(1) XML PRESERVE
WHITESPACE)
DSNU350I    -DBOB 300 20:39:13.13 DSNURRST - EXISTING RECORDS DELETED FROM TABLESPACE
DSNU304I    -DBOB 300 20:39:13.26 DSNURWT - (RE)LOAD PHASE STATISTICS - NUMBER OF RECORDS=1
FOR TABLE XMLR4.BK_TO_CSTMR_STMT
DSNU1147I   -DBOB 300 20:39:13.26 DSNURWT - (RE)LOAD PHASE STATISTICS - TOTAL NUMBER OF
RECORDS LOADED=1 FOR TABLESPACE DSN00242.BKRTORCS
DSNU302I    300 20:39:13.27 DSNURILD - (RE)LOAD PHASE STATISTICS - NUMBER OF INPUT RECORDS
PROCESSED=1
DSNU300I    300 20:39:13.27 DSNURILD - (RE)LOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU349I   -DBOB 300 20:39:13.32 DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=1 FOR
INDEX XMLR4.I_DOCIDBK_TO_CSTMR_STMT
DSNU258I    300 20:39:13.32 DSNURBXD - BUILD PHASE STATISTICS - NUMBER OF INDEXES=1
DSNU259I    300 20:39:13.32 DSNURBXD - BUILD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU010I    300 20:39:13.33 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0

```

NOTE: SYSREC shows the name of the data set which has the XML document. The first few characters of the XML document are shown below:

```

Û<?xml version="1.0" encoding="UTF-8" ?> <Document xmlns:xsi="http://ww ...
1F46A994A89A89977F4F748989889877EEC6F746644C98A989A4A999A7AA8778AA9766AA ...
1DCF74305592965EF1B0F055364957EF43608F0FE0C46344553074352A729EF8337A1166 ...

```

The first two positions contain X'11FD' which is the length of the XML document.

This technique requires specifying the exact size of the XML document preceding the data.

DB2 automatically generates the document ID column for each row that is loaded into a table with at least one XML column. The document ID column is partially hidden. It is not included in the result set of a SELECT * statement. However, you can query this column by name and view information about this column and its index in the catalog. Several utilities report information on this column in their output.

Loading XML data with the LOAD utility has the following restrictions:

- ▶ You cannot specify that XML input fields be loaded into non-XML columns, such as CHAR or VARCHAR columns.
- ▶ DB2 ignores any specified FREEPAGE and PCTFREE values until the next time that you run the REORG utility on this data.
- ▶ If you specify PREFORMAT, DB2 preformats the base table space, but not the XML table space.
- ▶ You cannot directly load the document ID column of the base table space.
- ▶ You cannot specify a default value for an XML column.
- ▶ You can load XML values that are greater than 32 KB by using file reference variables in the LOAD utility, or using applications with SQL XML as file reference variables.

LOAD utility using file reference variable

The method of loading XML records using file reference variables is used when the XML records are stored in separate input files. The normal input file contains the data for the non-XML columns of the base table and the names of the XML files. The sum of the length of all normal data fields and the XML file names cannot exceed 32 KB.

The XML input files can be any of these types:

- ▶ A sequential file
- ▶ A member of a PDS or PDSE
- ▶ A HFS file on a HFS directory
- ▶ A spanned file

The XML input file contains the entire XML record and the name of this file is stored in the normal load input file as a CHAR or VARCHAR field. So, instead of containing the whole XML record, the normal input file now only contains a file name, which in most cases no longer causes the sequential file to hit the 32 KB limit.

Additional keywords have been added to the CHAR and VARCHAR field specifications of the LOAD utility to support a file name as the input for the actual XML record:

- ▶ BLOBF: The input field contains the name of a file with a BLOB value.
- ▶ CLOBF: The input field contains the name of a file with a CLOB value.
- ▶ DBCLOBF: The input field contains the name of a file with a DBCLOB value.

In case of CLOBF and DBCLOBF, CCSID conversions are done when the CCSID of the input data is different than the CCSID of the table space. (EBCDIC, ASCII, UNICODE, or CCSID keywords might have been specified for the input data; the default is EBCDIC input data). In case of BLOBF, no conversions are done.

When the input field of a BLOBF, CLOBF, or DBCLOBF is NULL, the resulting XML value is NULL (null indicator field for the CHAR or VARCHAR field specified in the NULLIF keyword s hex FF).

Example 9-10 shows the JCL for the LOAD utility using file reference variable and the output of the utility run.

Example 9-10 LOAD utility JCL (using file reference variable) and output

```
//XMLR4LD JOB (999,POK),'DBOB COBOL',CLASS=A,
// MSGCLASS=T,NOTIFY=&SYSUID,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
/*
//LOAD1 EXEC DSNUPROC,SYSTEM=DBOB,UID=''
//SYSREC DD DSN=XMLR4.CVXML.XMLTEXT,DISP=SHR
//SYSERR DD DSN=XMLR4.LOAD.SYSERR,
// DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//SYSDISC DD DSN=XMLR4.LOAD.SYSDISC,
// DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//SYSMAP DD DSN=XMLR4.LOAD.SYSMAP,
// DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//SYSUT1 DD DSN=XMLR4.LOAD.SYSUT1,
// DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//UTPRINT DD SYSOUT=*
//SORTOUT DD DSN=XMLR4.LOAD.SORTOUT,
// DISP=(MOD,DELETE,CATLG),UNIT=SYSDA,
// SPACE=(4096,(20,20),,,ROUND)
//DSNUPROC.SYSIN DD *
LOAD DATA REPLACE
      INTO TABLE XMLR4.BK_TO_CSTMRT_STMT
      (BK_TO_CSTMRT_STMT POSITION(1:25) CHAR CLOBF)
/*
```

```

1DSNU000I    300 21:34:55.57 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = XMLR4.XMLR4LD
DSNU1044I    300 21:34:55.61 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I    300 21:34:55.62 DSNUGUTC - LOAD DATA REPLACE
DSNU650I    -DBOB 300 21:34:55.62 DSNURWI - INTO TABLE XMLR4.BK_TO_CSTMR_STMT
DSNU650I    -DBOB 300 21:34:55.62 DSNURWI - (BK_TO_CSTMR_STMT POSITION(1:25) CHAR CLOBF)
DSNU350I    -DBOB 300 21:34:56.31 DSNURRST - EXISTING RECORDS DELETED FROM TABLESPACE
DSNU304I    -DBOB 300 21:34:56.47 DSNURWT - (RE)LOAD PHASE STATISTICS - NUMBER OF RECORDS=1
FOR TABLE XMLR4.BK_TO_CSTMR_STMT
DSNU1147I   -DBOB 300 21:34:56.47 DSNURWT - (RE)LOAD PHASE STATISTICS - TOTAL NUMBER OF
RECORDS LOADED=1 FOR TABLESPACE DSN00242.BKRTORCS
DSNU302I    300 21:34:56.48 DSNURILD - (RE)LOAD PHASE STATISTICS - NUMBER OF INPUT RECORDS
PROCESSED=1
DSNU300I    300 21:34:56.48 DSNURILD - (RE)LOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU349I   -DBOB 300 21:34:56.54 DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=1 FOR
INDEX XMLR4.I_DOCIDBK_TO_CSTMR_STMT
DSNU258I    300 21:34:56.54 DSNURBXD - BUILD PHASE STATISTICS - NUMBER OF INDEXES=1
DSNU259I    300 21:34:56.54 DSNURBXD - BUILD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU010I    300 21:34:56.55 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0

```

NOTE: SYSREC shows the name of the data set where positions 1 to 25 have the name of the data set which has the XML document as shown below:

XMLR4.BKSTMT.XMLDATA1

The first few characters of the XML document are shown below:

```

<?xml version="1.0" encoding="UTF-8" ?> <Document xmlns:xsi="http://ww ...
46A994A89A89977F4F748989889877EEC6F746644C98A989A4A999A7AA8778AA9766AA ...
CF74305592965EF1B0F055364957EF43608F0FE0C46344553074352A729EF8337A1166 ...

```

The XML document starts from position 1.

LOAD XML data using input in binary format

Example 9-11 shows the JCL for the LOAD utility and the output of the utility run.

Example 9-11 LOAD utility JCL and output (input to load is in binary format)

```

//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
//*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SORTWK01 DD DSN=XMLR4.SORTWK01,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK02 DD DSN=XMLR4.SORTWK02,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK03 DD DSN=XMLR4.SORTWK03,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK04 DD DSN=XMLR4.SORTWK04,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSREC DD DSN=XMLR4.UNLOAD.XML.BINARY.RECOO,
// DISP=OLD

```

```

//DSNUPROC.SYSUT1 DD DSN=XMLR4.SYSUT1,
//  DISP=(MOD,DELETE,CATLG),
//  SPACE=(16384,(20,20),,,ROUND),
//  UNIT=SYSDA
//DSNUPROC.SORTOUT DD DSN=XMLR4.SORTOUT,
//  DISP=(MOD,DELETE,CATLG),
//  SPACE=(16384,(20,20),,,ROUND),
//  UNIT=SYSDA
//DSNUPROC.SYSERR DD DSN=XMLR4.SYSERR,
//  DISP=(MOD,DELETE,CATLG),
//  SPACE=(16384,(20,20),,,ROUND),
//  UNIT=SYSDA
//DSNUPROC.SYSMAP DD DSN=XMLR4.SYSMAP,
//  DISP=(MOD,DELETE,CATLG),
//  SPACE=(16384,(20,20),,,ROUND),
//  UNIT=SYSDA
//DSNUPROC.SYSIN DD *
LOAD DATA INDDN SYSREC LOG NO RESUME YES
EBCDIC CCSID(00037,00000,00000)
INTO TABLE "XMLR4"."BK_TO_CSTMRT_STMT"
WHEN(00001:00002) = X'0003'
NUMRECS 1
( "BK_TO_CSTMRT_STMT"
POSITION( 3) XML PRESERVE WHITESPACE BINARYXML)
DSNU000I 308 14:24:26.48 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 308 14:24:26.50 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 308 14:24:26.51 DSNUGUTC - LOAD DATA INDDN SYSREC LOG NO RESUME YES
EBCDIC CCSID(37, 0, 0)
DSNU650I -DBOB 308 14:24:26.51 DSNURWI - INTO TABLE "XMLR4"."BK_TO_CSTMRT_STMT"
WHEN(1:2)=X'0003' NUMRECS 1
DSNU650I -DBOB 308 14:24:26.51 DSNURWI - ("BK_TO_CSTMRT_STMT" POSITION(3) XML
PRESERVE WHITESPACE BINARYXML)
DSNU304I -DBOB 308 14:24:26.70 DSNURWT - (RE)LOAD PHASE STATISTICS - NUMBER OF
RECORDS=1 FOR TABLE XMLR4.BK_TO_CSTMRT_STMT
DSNU1147I -DBOB 308 14:24:26.70 DSNURWT - (RE)LOAD PHASE STATISTICS - TOTAL
NUMBER OF RECORDS LOADED=1 FOR TABLESPACE DSN00242.BKRTORCS
DSNU302I 308 14:24:26.70 DSNURILD - (RE)LOAD PHASE STATISTICS - NUMBER OF
INPUT RECORDS PROCESSED=1
DSNU300I 308 14:24:26.70 DSNURILD - (RE)LOAD PHASE COMPLETE, ELAPSED
TIME=00:00:00
DSNU349I -DBOB 308 14:24:26.76 DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF
KEYS=1 FOR INDEX XMLR4.I_DOCIDBK_TO_CSTMRT_STMT
DSNU258I 308 14:24:26.76 DSNURBXD - BUILD PHASE STATISTICS - NUMBER OF
INDEXES=1
DSNU259I 308 14:24:26.76 DSNURBXD - BUILD PHASE COMPLETE, ELAPSED
TIME=00:00:00
DSNU380I -DBOB 308 14:24:26.76 DSNUGRSX - TABLESPACE DSN00242.BKRTORCS PARTITION
1 IS IN COPY PENDING
DSNU380I -DBOB 308 14:24:26.76 DSNUGRSX - TABLESPACE DSN00242.XBKRO000 PARTITION
1 IS IN COPY PENDING
DSNU010I 308 14:24:26.77 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN
CODE=4

```

The LOAD control statement is the output of the UNLOAD utility run shown in Example 9-26 on page 223. We set POSITION(3) because positions 3 and 4 contain the length of the XML document. Apply the PTF for APAR PM29986 if you want to use POSITION(*).

LOAD XML data using input in spanned record format

Example 9-12 shows the JCL for the LOAD utility and the output of the utility run.

Example 9-12 LOAD utility JCL and output (input to load is in spanned record format)

```
//XMLR4LD JOB (999,POK), 'DBOB', CLASS=A,
// MSGCLASS=T, NOTIFY=XMLR4, TIME=1440, REGION=OM
/*JOBPARM SYSAFF=SC63, L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC, SYSTEM=DBOB, UID='TEMP', UTPROC=' '
//DSNUPROC.SORTWK01 DD DSN=XMLR4.SORTWK01,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK02 DD DSN=XMLR4.SORTWK02,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK03 DD DSN=XMLR4.SORTWK03,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK04 DD DSN=XMLR4.SORTWK04,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSREC DD DSN=XMLR4.DSN00242.XBKRO000.T214620.UFILEREF.NEW,
// DISP=OLD
//DSNUPROC.SYSUT1 DD DSN=XMLR4.SYSUT1,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTOUT DD DSN=XMLR4.SORTOUT,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSERR DD DSN=XMLR4.SYSERR,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSMAP DD DSN=XMLR4.SYSMAP,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSIN DD *
LOAD DATA INDDN SYSREC LOG NO RESUME YES
EBCDIC CCSID(00037,00000,00000)
FORMAT SPANNED YES
INTO TABLE "XMLR4"."BK_TO_CSTMRT_STMT"
WHEN(00001:00002) = X'0003'
NUMRECS 1
```

```
( "DSN_NULL_IND_00001" POSITION( 00003)          CHAR(1)
, "MSG_ID"
  POSITION( 00004)          VARCHAR
                        NULLIF(DSN_NULL_IND_00001)=X'FF'
, "DSN_NULL_IND_00002" POSITION( *)              CHAR(1)
, "MSG_CRE_DT_TM"
  POSITION( *)              TIMESTAMP WITH TIME ZONE EXTERNAL(032)
                        NULLIF(DSN_NULL_IND_00002)=X'FF'
, "BK_TO_CSTMRTM"
  POSITION( *)              XML PRESERVE WHITESPACE )
```

NOTE: The LOAD control statement is the output of UNLOAD utility run shown in Example 9-27 on page 224.

```
1DSNU000I 309 19:07:40.66 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 309 19:07:40.69 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 309 19:07:40.70 DSNUGUTC - LOAD DATA INDDN SYSREC LOG NO RESUME YES EBCDIC CCSID(37, 0, 0) FORMAT SPANNED YES
DSNU650I -DBOB 309 19:07:40.70 DSNURWI - INTO TABLE "XMLR4". "BK_TO_CSTMRTM" WHEN(1:2)=X'0003' NUMRECS 1
DSNU650I -DBOB 309 19:07:40.70 DSNURWI - ("DSN_NULL_IND_00001" POSITION(3) CHAR(1),
DSNU650I -DBOB 309 19:07:40.70 DSNURWI - "MSG_ID" POSITION(4) VARCHAR NULLIF(DSN_NULL_IND_00001)=X'FF',
DSNU650I -DBOB 309 19:07:40.70 DSNURWI - "DSN_NULL_IND_00002" POSITION(*) CHAR(1),
DSNU650I -DBOB 309 19:07:40.70 DSNURWI - "MSG_CRE_DT_TM" POSITION(*) TIMESTAMP WITH TIME ZONE EXTERNAL(32)
NULLIF(DSN_NULL_IND_00002)=X'FF',
DSNU650I -DBOB 309 19:07:40.70 DSNURWI - "BK_TO_CSTMRTM" POSITION(*) XML PRESERVE WHITESPACE)
DSNU304I -DBOB 309 19:07:40.89 DSNURWT - (RE)LOAD PHASE STATISTICS - NUMBER OF RECORDS=1 FOR TABLE XMLR4.BK_TO_CSTMRTM
DSNU1147I -DBOB 309 19:07:40.89 DSNURWT - (RE)LOAD PHASE STATISTICS - TOTAL NUMBER OF RECORDS LOADED=1 FOR TABLESPACE DSN00242.BKRTORCS
DSNU302I 309 19:07:40.89 DSNURILD - (RE)LOAD PHASE STATISTICS - NUMBER OF INPUT RECORDS PROCESSED=1
DSNU300I 309 19:07:40.89 DSNURILD - (RE)LOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU349I -DBOB 309 19:07:40.95 DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=1 FOR INDEX XMLR4.I_DOCIDBK_TO_CSTMRTM
DSNU258I 309 19:07:40.95 DSNURBXD - BUILD PHASE STATISTICS - NUMBER OF INDEXES=1
DSNU259I 309 19:07:40.95 DSNURBXD - BUILD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU380I -DBOB 309 19:07:40.96 DSNUGSRX - TABLESPACE DSN00242.BKRTORCS PARTITION 1 IS IN COPY PENDING
DSNU380I -DBOB 309 19:07:40.96 DSNUGSRX - TABLESPACE DSN00242.XBKRO000 PARTITION 1 IS IN COPY PENDING
DSNU010I 309 19:07:40.96 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=4
```

For UNLOAD, SPANNED with BINARY XML performs significantly better than other formats. Use SPANNED if possible rather than file references.

9.8 MERGECOPY

MERGECOPY utility merges incremental image copies to produce a merged incremental image copy, or merge full image copy with incremental image copies to produce a full image copy of the XML table spaces.

Example 9-13 shows the JCL for MERGECOPY utility and the output of the utility run. This utility run merges the full image copy taken in Example 9-3 on page 188 and the incremental image copy taken in Example 9-4 on page 190.

Example 9-13 MERGECOPY utility JCL and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=OM
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
TEMPLATE A DSN(&DB..&SN..&IC..D&DATE..T&TIME..COPY)
MERGECOPY TABLESPACE DSN00242.BKRTORCS COPYDDN(A) NEWCOPY YES
MERGECOPY TABLESPACE DSN00242.XBKRO000 COPYDDN(A) NEWCOPY YES
```

```
1DSNU000I 314 18:30:00.98 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 314 18:30:01.01 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 314 18:30:01.02 DSNUGUTC - TEMPLATE A
DSN(&DB..&SN..&IC..D&DATE..T&TIME..COPY)
DSNU1035I 314 18:30:01.02 DSNUJTRD - TEMPLATE STATEMENT PROCESSED SUCCESSFULLY
```

```

ODSNU050I 314 18:30:01.02 DSNUGUTC - MERGECOPY TABLESPACE DSN00242.BKRTORCS COPYDDN(A)
NEWCOPY YES
DSNU1038I 314 18:30:01.06 DSNUGDYN - DATASET ALLOCATED. TEMPLATE=A
          DDNAME=SYS00001
          DSN=DSN00242.BKRTORCS.F.D2010314.T233001.COPY
DSNU463I 314 18:30:01.07 DSNUYBR3 - THE PRIMARY IMAGE COPY DATA SET
DSN00242.BKRTORCS.I.D2010314.T231425.COPY WITH DATE=101110 AND TIME=181425
          IS PARTICIPATING IN MERGECOPY.
DSNU463I 314 18:30:01.09 DSNUYBR3 - THE PRIMARY IMAGE COPY DATA SET
DSN00242.BKRTORCS.F.D2010314.T231248.COPY WITH DATE=101110 AND TIME=181248
          IS PARTICIPATING IN MERGECOPY.
DSNU454I 314 18:30:01.13 DSNUYBRO - COPY MERGE COMPLETE
          NUMBER OF COPIES=2
          NUMBER OF COPIES MERGED=2
          TOTAL NUMBER OF PAGES MERGED=6
          ELAPSED TIME=00:00:00

ODSNU050I 314 18:30:01.14 DSNUGUTC - MERGECOPY TABLESPACE DSN00242.XBKRO000 COPYDDN(A)
NEWCOPY YES
DSNU1038I 314 18:30:01.17 DSNUGDYN - DATASET ALLOCATED. TEMPLATE=A
          DDNAME=SYS00004
          DSN=DSN00242.XBKRO000.F.D2010314.T233001.COPY
DSNU463I 314 18:30:01.18 DSNUYBR3 - THE PRIMARY IMAGE COPY DATA SET
DSN00242.XBKRO000.I.D2010314.T231425.COPY WITH DATE=101110 AND TIME=181425
          IS PARTICIPATING IN MERGECOPY.
DSNU463I 314 18:30:01.19 DSNUYBR3 - THE PRIMARY IMAGE COPY DATA SET
DSN00242.XBKRO000.F.D2010314.T231248.COPY WITH DATE=101110 AND TIME=181248
          IS PARTICIPATING IN MERGECOPY.
DSNU454I 314 18:30:01.23 DSNUYBRO - COPY MERGE COMPLETE
          NUMBER OF COPIES=2
          NUMBER OF COPIES MERGED=2
          TOTAL NUMBER OF PAGES MERGED=6
          ELAPSED TIME=00:00:00

```

9.9 QUIESCE

When you specify QUIESCE TABLESPACESET, the table space set includes related XML objects. You can specify that you want to quiesce the XML table space, base table space, or both of them, plus related index spaces if they are copy enabled.

All table spaces that are involved in a versioning relationship are quiesced when QUIESCE is run on either the system-maintained temporal table or the history table space. Auxiliary LOB and XML table spaces on both system-maintained temporal table spaces and history table spaces are included.

Example 9-14 shows the JCL for the QUIESCE utility and the output of the utility run.

Example 9-14 QUIESCE utility JCL and output (1 of 2)

```

//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
QUIESCE TABLESPACESET DSN00242.BKRTORCS

```

```

1DSNU000I   314 21:50:34.75 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
  DSNU1044I   314 21:50:34.77 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I   314 21:50:34.78 DSNUGUTC - QUIESCE TABLESPACESET DSN00242.BKRTORCS
  DSNU477I   -DBOB 314 21:50:34.78 DSNUQUIA - QUIESCE SUCCESSFUL FOR TABLESPACESET
  DSN00242.BKRTORCS
  DSNU477I   -DBOB 314 21:50:34.78 DSNUQUIA - QUIESCE SUCCESSFUL FOR TABLESPACE
  DSN00242.BKRTORCS
  DSNU477I   -DBOB 314 21:50:34.78 DSNUQUIA - QUIESCE SUCCESSFUL FOR TABLESPACE
  DSN00242.XBKRO000
  DSNU474I   -DBOB 314 21:50:34.78 DSNUQUIA - QUIESCE AT RBA 000069A77034 AND AT LRSN
  000069A77034
  DSNU475I   314 21:50:34.78 DSNUQUIB - QUIESCE UTILITY COMPLETE, ELAPSED TIME= 00:00:00
  DSNU010I   314 21:50:34.79 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0

```

Notice that only the table spaces are included. Indexes are included only if they are copy enabled.

Let us alter the index definitions to make them copy enabled. Figure 9-6 shows how this is done.

```

SELECT NAME,DBNAME,COPY FROM SYSIBM.SYSINDEXES
WHERE DBNAME='DSN00242';
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
NAME                                         DBNAME                                         COPY
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
I_NODEIDXBK_TO_CSTMR_STMT                   DSN00242                                       N
I_DOCIDBK_TO_CSTMR_STMT                     DSN00242                                       N
DSNE610I NUMBER OF ROWS DISPLAYED IS 2
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
ALTER INDEX I_NODEIDXBK_TO_CSTMR_STMT COPY YES;
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 0
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
ALTER INDEX I_DOCIDBK_TO_CSTMR_STMT COPY YES;
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 0
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
SELECT NAME,DBNAME,COPY FROM SYSIBM.SYSINDEXES
WHERE DBNAME='DSN00242';
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
NAME                                         DBNAME                                         COPY
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
I_NODEIDXBK_TO_CSTMR_STMT                   DSN00242                                       Y
I_DOCIDBK_TO_CSTMR_STMT                     DSN00242                                       Y
DSNE610I NUMBER OF ROWS DISPLAYED IS 2
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100

```

Figure 9-6 Copy enable the DOCID and NODEID indexes

Example 9-15 shows the JCL for the QUIESCE utility and the output of the utility run.

Example 9-15 QUIESCE utility JCL and output (2 of 2)

```

//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999

```

```
// JCLLIB ORDER=(DB0BM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DB0B,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
QUIESCE TABLESPACESET DSN00242.BKRTORCS
1DSNU000I 314 22:14:39.70 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 314 22:14:39.73 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
0DSNU050I 314 22:14:39.73 DSNUGUTC - QUIESCE TABLESPACESET DSN00242.BKRTORCS
DSNU477I -DB0B 314 22:14:39.74 DSNUQUIA - QUIESCE SUCCESSFUL FOR TABLESPACESET
DSN00242.BKRTORCS
DSNU477I -DB0B 314 22:14:39.74 DSNUQUIA - QUIESCE SUCCESSFUL FOR TABLESPACE
DSN00242.BKRTORCS
DSNU477I -DB0B 314 22:14:39.74 DSNUQUIA - QUIESCE SUCCESSFUL FOR INDEXSPACE
DSN00242.IRDOCIDB
DSNU477I -DB0B 314 22:14:39.74 DSNUQUIA - QUIESCE SUCCESSFUL FOR TABLESPACE
DSN00242.XBKRO000
DSNU477I -DB0B 314 22:14:39.74 DSNUQUIA - QUIESCE SUCCESSFUL FOR INDEXSPACE
DSN00242.IRNODEID
DSNU474I -DB0B 314 22:14:39.74 DSNUQUIA - QUIESCE AT RBA 000069A99034 AND AT LRSN
000069A99034
DSNU568I -DB0B 314 22:14:39.74 DSNUGSRX - INDEX XMLR4.I_DOCIDBK_TO_CSTMRT_STMT IS IN
INFORMATIONAL COPY PENDING STATE
DSNU568I -DB0B 314 22:14:39.74 DSNUGSRX - INDEX XMLR4.I_NODEIDXBK_TO_CSTMRT_STMT IS IN
INFORMATIONAL COPY PENDING STATE
DSNU475I 314 22:14:39.74 DSNUQUIB - QUIESCE UTILITY COMPLETE, ELAPSED TIME= 00:00:00
```

As expected, indexes are now included.

9.10 REBUILD INDEX

You can use the REBUILD INDEX utility to rebuild XML indexes, DOCID indexes, and NODEID indexes. You do not need to specify any additional keywords in the REBUILD INDEX statement. SHRLEVEL CHANGE is not allowed on not logged tables and XML indexes.

When you process both NODEID indexes and XML indexes together, they are processed sequentially. First the NODEID index is processed and then the XML index.

Example 9-16 shows the JCL for the REBUILD utility and the output of the utility run

Example 9-16 REBUILD INDEX utility JCL and output

```
//XMLR4LD JOB (999,POK),'DB0B',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=OM
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DB0BM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DB0B,UID='TEMP',UTPROC=''
//DSNUPROC.SORTWK01 DD DSN=XMLR4.SORTWK01,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK02 DD DSN=XMLR4.SORTWK02,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK03 DD DSN=XMLR4.SORTWK03,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK04 DD DSN=XMLR4.SORTWK04,
```

```

// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSUT1 DD DSN=XMLR4.SYSUT1,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSIN DD *
REBUILD INDEX (ALL) TABLESPACE DSN00242.BKRTORCS
REBUILD INDEX (ALL) TABLESPACE DSN00242.XBKRO000
1DSNU000I 315 19:54:04.76 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 315 19:54:04.79 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 315 19:54:04.79 DSNUGUTC - REBUILD INDEX(ALL) TABLESPACE DSN00242.BKRTORCS
DSNU3343I -DBOB 315 19:54:04.80 DSNUCINM - REAL-TIME STATISTICS INFORMATION MISSING FOR
TABLESPACE DSN00242.BKRTORCS PARTITION 1
DSNU555I -DBOB 315 19:54:04.81 DSNUCRUL - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
PROCESSED=1
DSNU705I 315 19:54:04.81 DSNUCRIB - UNLOAD PHASE COMPLETE - ELAPSED TIME=00:00:00
DSNU394I -DBOB 315 19:54:05.01 DSNURBXC - SORTBLD PHASE STATISTICS - NUMBER OF KEYS=1 FOR
INDEX XMLR4.I_DOCIDBK_TO_CSTMRT_STMT
DSNU391I 315 19:54:05.01 DSNUCRIB - SORTBLD PHASE STATISTICS. NUMBER OF INDEXES = 1
DSNU392I 315 19:54:05.01 DSNUCRIB - SORTBLD PHASE COMPLETE, ELAPSED TIME = 00:00:00
DSNU425I -DBOB 315 19:54:05.03 DSNUGFCR - INDEXSPACE DSN00242.IRDOCIDB DOES NOT HAVE
THE COPY YES ATTRIBUTE
ODSNU050I 315 19:54:05.04 DSNUGUTC - REBUILD INDEX(ALL) TABLESPACE DSN00242.XBKRO000
DSNU3343I -DBOB 315 19:54:05.05 DSNUCINM - REAL-TIME STATISTICS INFORMATION MISSING FOR
TABLESPACE DSN00242.XBKRO000 PARTITION 1
DSNU555I -DBOB 315 19:54:05.06 DSNUCRUL - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
PROCESSED=1
DSNU705I 315 19:54:05.06 DSNUCRIB - UNLOAD PHASE COMPLETE - ELAPSED TIME=00:00:00
DSNU394I -DBOB 315 19:54:05.27 DSNURBXC - SORTBLD PHASE STATISTICS - NUMBER OF KEYS=1 FOR
INDEX XMLR4.I_NODEIDXBK_TO_CSTMRT_STMT
DSNU391I 315 19:54:05.27 DSNUCRIB - SORTBLD PHASE STATISTICS. NUMBER OF INDEXES = 1
DSNU392I 315 19:54:05.27 DSNUCRIB - SORTBLD PHASE COMPLETE, ELAPSED TIME = 00:00:00
DSNU425I -DBOB 315 19:54:05.28 DSNUGFCR - INDEXSPACE DSN00242.IRNODEID DOES NOT HAVE
THE COPY YES ATTRIBUTE
DSNU010I 315 19:54:05.28 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0

```

9.11 RECOVER INDEX and RECOVER TABLESPACE

You can use the RECOVER utility to recover XML objects. You do not need to specify any additional keywords in the RECOVER statement. When you recover an XML table space or index to a point in time, you should recover all related objects to the same point in time. Related objects include XML objects, LOB objects, and referentially related objects. If you do not recover all related objects to the same point in time, one or more objects might be placed in a restrictive state.

For point-in-time recoveries of base, LOB, XML, and history objects, the option VERIFYSET, new with DB2 10, specifies whether the RECOVER utility verifies that all related objects that are required for the point-in-time recovery are included in the RECOVER control statement.

- *VERIFYSET YES* means the RECOVER utility verifies that all of the objects that are required to perform a point-in-time recovery of the base, LOB, XML, and history objects, have been included in the RECOVER control statement. VERIFYSET YES is the default.

- ▶ *VERIFYSET NO* means the RECOVER utility does not verify that all of the objects that are required to perform a point-in-time recovery of the base, LOB, XML, and history objects, have been included in the RECOVER control statement.

Specifying VERIFYSET NO allows you to break up a point-in-time recovery into multiple jobs or to avoid recovering objects that have changed since the selected recovery point.

The VERIFYSET option does not apply to point-in-time recoveries of catalog and directory objects.

The other option for point-in-time recovery is ENFORCE YES/NO.

- ▶ *ENFORCE YES* specifies that CHKP and ACHKP pending states are set for a point-in-time recovery when only a subset of the related objects (BASE, LOB, XML, and RI) have been recovered to a point in time. ENFORCE YES is the default for catalog and directory objects. There is no override for the ENFORCE YES option for catalog and directory objects.
- ▶ *ENFORCE NO* specifies that CHKP and ACHKP pending states are not set for a point-in-time recovery when only a subset of the related objects (BASE, LOB, XML, and RI) have been recovered to a point in time.

We took a full image copy of the base table space and XML table space (See Example 9-3 on page 188), did a partial update to the XML document (See Example on page 190), took an incremental image copy (See Example 9-4 on page 190), and merged the full and incremental image copies (See Example 9-13 on page 202).

We now want to take the table space back to the full image copy before we dis the partial update to the XML document.

Before we do, we examine the status of the database and query the content of the table.

Figure 9-7 shows the status of the database.

```

-DISPLAY DB(DSN00242)

DSNT360I  -DBOB *****
DSNT361I  -DBOB *  DISPLAY DATABASE SUMMARY
           *  GLOBAL
DSNT360I  -DBOB *****
DSNT362I  -DBOB      DATABASE = DSN00242  STATUS = RW
           DBD LENGTH = 4028

DSNT397I  -DBOB
NAME      TYPE PART  STATUS          PHYERRLO PHYERRHI CATALOG  PIECE
-----
BKRTORCS TS      0001 RW
BKRTORCS TS              RW
XBKR0000 XS      0001 RW
XBKR0000 XS              RW
IRDOCIDB IX      L0001 RW
IRDOCIDB IX        L*  RW
IRNODEID IX      L0001 RW
IRNODEID IX        L*  RW
*****  DISPLAY OF DATABASE DSN00242 ENDED  *****
DSN9022I  -DBOB DSNTDDIS 'DISPLAY DATABASE' NORMAL COMPLETION
***

```

Figure 9-7 Status of database DSN00242

Figure 9-8 shows the content of the XML document.

```
SELECT XMLSERIALIZE(
      XMLQUERY(
        'declare default element namespace
         "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
         /Document/BkToCstmrStmt/Stmt/Bal/Amt[../Tp/CdOrPrtry/Cd="CLBD"] '
        PASSING BK_TO_CSTMRTMT) AS CLOB(500))
FROM BK_TO_CSTMRTMT;

<Amt xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02" Ccy="SEK">900000</Amt>
DSNE610I NUMBER OF ROWS DISPLAYED IS 1
```

Figure 9-8 Content of the XML document

Example 9-17 shows the JCL for RECOVER utility to recover the base table space to the LRSN value associated with the full image copy taken before the partial update to the XML document was done.

Example 9-17 RECOVER TABLESPACE utility JCL and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
RECOVER TABLESPACE DSN00242.BKRTORCS TOLOGPOINT X'000069667C5E'
1DSNU000I 315 19:49:49.27 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 315 19:49:49.30 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 315 19:49:49.30 DSNUGUTC - RECOVER TABLESPACE DSN00242.BKRTORCS TOLOGPOINT
X'000011112222'
DSNU1316I -DBOB 315 19:49:49.31 DSNUCAIN - THE FOLLOWING TABLESPACES ARE MISSING FROM THE
RECOVERY LIST DSN00242.XBKRO000
DSNU500I 315 19:49:49.31 DSNUCBDR - RECOVERY COMPLETE, ELAPSED TIME=00:00:00
DSNU012I 315 19:49:49.31 DSNUGBAC - UTILITY EXECUTION TERMINATED, HIGHEST RETURN CODE=8
```

Note: The RBA value for the TOLOGPOINT is the LRSN value associated with the full image copy shown in Example 9-21 on page 216. It is important to ensure that both the base table space and the XML table space are recovered to the same point-in-time.

Example 9-18 shows the JCL with the revised RECOVER utility control statement and output of the utility run.

Example 9-18 RECOVER TABLESPACE utility JCL (with modified control statement) and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
RECOVER TABLESPACE DSN00242.BKRTORCS
TABLESPACE DSN00242.XBKRO000
TOLOGPOINT X'000069667C5E'
```



```

1DSNU000I    315 19:36:29.07 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
  DSNU1044I    315 19:36:29.10 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I    315 19:36:29.10 DSNUGUTC - RECOVER TABLESPACE DSN00242.BKRTORCS TABLESPACE
DSN00242.XBKRO000
  TOLOGPOINT X'000069667C5E'
DSNU532I    315 19:36:29.11 DSNUCBMD - RECOVER TABLESPACE DSN00242.BKRTORCS  START
DSNU515I    315 19:36:29.11 DSNUCBAL - THE IMAGE COPY DATA SET
DSN00242.BKRTORCS.F.D2010314.T231248.COPY WITH
  DATE=20101110 AND TIME=181248
  IS PARTICIPATING IN RECOVERY OF TABLESPACE DSN00242.BKRTORCS
DSNU504I    315 19:36:29.33 DSNUCBMD - MERGE STATISTICS FOR TABLESPACE DSN00242.BKRTORCS
-
  NUMBER OF COPIES=1
  NUMBER OF PAGES MERGED=3
  ELAPSED TIME=00:00:00
DSNU532I    315 19:36:29.34 DSNUCBMD - RECOVER TABLESPACE DSN00242.XBKRO000  START
DSNU515I    315 19:36:29.34 DSNUCBAL - THE IMAGE COPY DATA SET
DSN00242.XBKRO000.F.D2010314.T231248.COPY WITH
  DATE=20101110 AND TIME=181248
  IS PARTICIPATING IN RECOVERY OF TABLESPACE DSN00242.XBKRO000
DSNU504I    315 19:36:29.54 DSNUCBMD - MERGE STATISTICS FOR TABLESPACE DSN00242.XBKRO000
-
  NUMBER OF COPIES=1
  NUMBER OF PAGES MERGED=3
  ELAPSED TIME=00:00:00
DSNU830I  -DBOB 315 19:36:29.14 DSNUCARS - INDEX XMLR4.I_DOCIDBK_TO_CSTMR_STMT IS IN
REBUILD PENDING
DSNU831I  -DBOB 315 19:36:29.14 DSNUCARS - ALL INDEXES OF DSN00242.BKRTORCS ARE IN REBUILD
PENDING
DSNU830I  -DBOB 315 19:36:29.36 DSNUCARS - INDEX XMLR4.I_NODEIDXBK_TO_CSTMR_STMT IS IN
REBUILD PENDING
DSNU831I  -DBOB 315 19:36:29.36 DSNUCARS - ALL INDEXES OF DSN00242.XBKRO000 ARE IN REBUILD
PENDING
DSNU1511I -DBOB 315 19:36:29.55 DSNUCALA - FAST LOG APPLY WAS NOT USED FOR RECOVERY
DSNU1510I    315 19:36:29.55 DSNUCBLA - LOG APPLY PHASE COMPLETE, ELAPSED TIME = 00:00:00
DSNU535I  -DBOB 315 19:36:29.55 DSNUCATM - FOLLOWING TABLESPACES RECOVERED TO A CONSISTENT
POINT
  DSN00242.BKRTORCS
  DSN00242.XBKRO000
DSNU500I    315 19:36:29.57 DSNUCBDR - RECOVERY COMPLETE, ELAPSED TIME=00:00:00

```

The base table space and XML table space are now recovered to the same point-in-time. DB2 places the DOCID and NODEID indexes in REBUILD-pending state (RBDP) to ensure the indexes are consistent with the data.

Figure 9-9 shows the status of the database.

```

-DISPLAY DB(DSN00242)

DSNT360I  -DBOB *****
DSNT361I  -DBOB *  DISPLAY DATABASE SUMMARY
           *    GLOBAL
DSNT360I  -DBOB *****
DSNT362I  -DBOB      DATABASE = DSN00242  STATUS = RW
           DBD LENGTH = 4028
DSNT397I  -DBOB
NAME      TYPE PART  STATUS              PHYERRLO PHYERRHI CATALOG  PIECE
-----
BKRTORCS TS      0001 RW
BKRTORCS TS              RW
XBKR0000 XS      0001 RW
XBKR0000 XS              RW
IRDOCIDB IX      L0001 RW RBDP
IRDOCIDB IX        L*  RW RBDP
IRNODEID IX      L0001 RW RBDP
IRNODEID IX        L*  RW RBDP
*****  DISPLAY OF DATABASE DSN00242 ENDED          *****
DSN9022I  -DBOB DSNTDDIS 'DISPLAY DATABASE' NORMAL COMPLETION
***

```

Figure 9-9 Status of database DSN00242 (after partial recovery)

Let us rebuild the indexes. Example 9-16 on page 205 shows the JCL for the REBUILD INDEX utility and the output of the utility run.

Figure 9-10 shows the status of the database after the indexes are rebuilt.

```

-DISPLAY DB(DSN00242)

DSNT360I  -DBOB *****
DSNT361I  -DBOB *  DISPLAY DATABASE SUMMARY
           *    GLOBAL
DSNT360I  -DBOB *****
DSNT362I  -DBOB      DATABASE = DSN00242  STATUS = RW
           DBD LENGTH = 4028
DSNT397I  -DBOB
NAME      TYPE PART  STATUS              PHYERRLO PHYERRHI CATALOG  PIECE
-----
BKRTORCS TS      0001 RW
BKRTORCS TS              RW
XBKR0000 XS      0001 RW
XBKR0000 XS              RW
IRDOCIDB IX      L0001 RW
IRDOCIDB IX        L*  RW
IRNODEID IX      L0001 RW
IRNODEID IX        L*  RW
*****  DISPLAY OF DATABASE DSN00242 ENDED          *****
DSN9022I  -DBOB DSNTDDIS 'DISPLAY DATABASE' NORMAL COMPLETION
***

```

Figure 9-10 Status of database DSN00242 (after partial recovery and rebuild of indexes)

Figure 9-11 shows the content of the XML document.

```
SELECT XMLSERIALIZE(
      XMLQUERY(
'declare default element namespace
"urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
  /Document/BkToCstmrStmt/Stmt/Bal/Amt[../Tp/CdOrPrtry/Cd="CLBD"] '
      PASSING BK_TO_CSTMRTMT) AS CLOB(500))
FROM BK_TO_CSTMRTMT;

<Amt xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02"
Ccy="SEK">435678.50</Amt>
DSNE610I NUMBER OF ROWS DISPLAYED IS 1
```

Figure 9-11 Content of the XML document after partial recovery

Note: If the DOCID and NODEID indexes are enabled for image copy, in Figure 9-9 on page 210 you would see RECP instead of RBDP. Instead of rebuilding the indexes you would recover the indexes using RECOVER INDEX utility which uses the image copy you would have established for these indexes using the COPY utility.

9.12 REORG INDEX and REORG TABLESPACE

You can use the REORG INDEX utility to reorganize XML indexes, and the REORG TABLESPACE utility to reorganize XML table space. You do not need to specify any additional keywords in the REORG statement. When you specify that you want XML objects (either XML indexes or XML table spaces) to be reorganized, you must also specify the WORKDDN keyword and provide the specified temporary work file. The default is SYSUT1.

When you specify the name of the base table space in the REORG statement, DB2 reorganizes only that table space and not any related XML objects. If you want DB2 to reorganize the XML objects, you must specify those object names.

When you run REORG on an XML table space that supports XML versions, REORG discards rows for versions of an XML document that are no longer needed.

For XML table spaces and base table spaces with XML columns, you cannot specify the following options in the REORG statement:

- ▶ DISCARD
- ▶ REBALANCE
- ▶ UNLOAD EXTERNAL

REORG can take inline copies of XML table spaces.

If you specify a base table space with the STATISTICS keyword, DB2 does not gather statistics for the related XML table space or its indexes.

If large amounts of data are deleted from a partition-by-growth universal table space, including XML table spaces, run the REORG TABLESPACE utility with SHRLEVEL

REFERENCE or SHRLEVEL CHANGE on the entire table space to reclaim physical space from the partition-by-growth and XML table spaces.

Do not use REORG UNLOAD ONLY to propagate data. When you specify the UNLOAD ONLY option, REORG unloads only the data that physically resides in the base table space; LOB and XML columns are not unloaded. For purposes of data propagation, you should use UNLOAD or REORG UNLOAD EXTERNAL instead.

Example 9-19 shows the JCL for the REORG TABLESPACE utility and the output of the utility run.

Example 9-19 REORG TABLESPACE utility JCL and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=OM
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SORTWK01 DD DSN=XMLR4.SORTWK01,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK02 DD DSN=XMLR4.SORTWK02,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK03 DD DSN=XMLR4.SORTWK03,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTWK04 DD DSN=XMLR4.SORTWK04,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSREC DD DSN=XMLR4.DSN00242.REORG.DATA,
// DISP=(MOD,CATLG)
//DSNUPROC.SYSUT1 DD DSN=XMLR4.SYSUT1,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SORTOUT DD DSN=XMLR4.SORTOUT,
// DISP=(MOD,DELETE,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSIN DD *
REORG TABLESPACE DSN00242.BKRTORCS
REORG TABLESPACE DSN00242.XBKR0000

1DSNU000I 319 18:06:00.66 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 319 18:06:00.68 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 319 18:06:00.69 DSNUGUTC - REORG TABLESPACE DSN00242.BKRTORCS
DSNU251I 319 18:06:00.81 DSNURULD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS UNLOADED=1 FOR TABLESPACE
DSN00242.BKRTORCS PART 1
DSNU252I 319 18:06:00.81 DSNURULD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS UNLOADED=1 FOR TABLESPACE
DSN00242.BKRTORCS
DSNU250I 319 18:06:00.81 DSNURULD - UNLOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU303I -DBOB 319 18:06:01.21 DSNURWT - (RE)LOAD PHASE STATISTICS - NUMBER OF RECORDS=1 FOR TABLE
XMLR4.BK TO CSTMR STMT PART=1
DSNU304I -DBOB 319 18:06:01.21 DSNURWT - (RE)LOAD PHASE STATISTICS - NUMBER OF RECORDS=1 FOR TABLE
XMLR4.BK TO CSTMR STMT
DSNU302I 319 18:06:01.22 DSNURILD - (RE)LOAD PHASE STATISTICS - NUMBER OF INPUT RECORDS PROCESSED=1
DSNU300I 319 18:06:01.22 DSNURILD - (RE)LOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU042I 319 18:06:01.22 DSNUGSOR - SORT PHASE STATISTICS -
NUMBER OF RECORDS=1
ELAPSED TIME=00:00:00
DSNU349I -DBOB 319 18:06:01.28 DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=1 FOR INDEX
XMLR4.I_DOCIDBK TO CSTMR STMT
DSNU258I 319 18:06:01.28 DSNURBXD - BUILD PHASE STATISTICS - NUMBER OF INDEXES=1
DSNU259I 319 18:06:01.28 DSNURBXD - BUILD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU421I 319 18:06:01.30 DSNUGFUM - START OF DFSMS MESSAGES

1PAGE 0001 5695-DF175 DFSMSDSS VIR12.0 DATA SET SERVICES 2010.319 18:06
-ADRO30I (SCH)-PRIME( 0), DCB VALUES HAVE BEEN MODIFIED FOR SYSPRINT. BLKSIZE VALUE MODIFIED FROM 0 TO 128
COPY DATASET(INCLUDE( -
DBOBD.DSNDBC.DSN00242.BKRTORCS.I0001.A001 )) -
RENAMEU( -
(DBOBD.DSNDBC.DSN00242.BKRTORCS.I0001.A001 , -
DBOBI.DSN00242.BKRTORCS.N00001.C2RPDUCT )) -
REPUNC ALLDATA(*) ALLEXCP CANCELERROR SHARE -
WRITECHECK TOLERATE(ENQF)
ADR101I (R/I)-RI01 (01), TASKID 001 HAS BEEN ASSIGNED TO COMMAND 'COPY '
ADR109I (R/I)-RI01 (01), 2010.319 18:06:01 INITIAL SCAN OF USER CONTROL STATEMENTS COMPLETED
ADRO50I (001)-PRIME(01), DFSMSDSS INVOKED VIA APPLICATION INTERFACE
ADRO16I (001)-PRIME(01), RACF® LOGGING OPTION IN EFFECT FOR THIS TASK
OADRO06I (001)-STEND(01), 2010.319 18:06:01 EXECUTION BEGINS
```

```

OADR711I (001)-NEWS(01), DATA SET DBOBD.DSNDBC.DSN00242.BKRTORCS.I0001.A001 HAS BEEN ALLOCATED WITH NEWNAME
      DBOBI.DSN00242.BKRTORCS.N00001.C2RPDUCT USING STORCLAS DBOBDATA, DATACLAS DBOB, AND
MGMTCLAS MCDB22
OADR806I (001)-TOMI (03), DATA SET DBOBD.DSNDBC.DSN00242.BKRTORCS.I0001.A001 COPIED USING A FAST REPLICATION
FUNCTION
OADR801I (001)-DDDS (01), DATA SET FILTERING IS COMPLETE. 1 OF 1 DATA SETS WERE SELECTED: 0 FAILED
SERIALIZATION AND 0 FAILED FOR OTHER REASONS
OADR454I (001)-DDDS (01), THE FOLLOWING DATA SETS WERE SUCCESSFULLY PROCESSED
0      DBOBD.DSNDBC.DSN00242.BKRTORCS.I0001.A001
OADR006I (001)-STEND(02), 2010.319 18:06:01 EXECUTION ENDS
OADR013I (001)-CLTSK(01), 2010.319 18:06:01 TASK COMPLETED WITH RETURN CODE 0000
OADR012I (SCH)-DSSU (01), 2010.319 18:06:01 DFSMSDSS PROCESSING COMPLETE. HIGHEST RETURN CODE IS 0000
DSNU422I 319 18:06:01.48 DSNUGFCD - END OF DFSMS MESSAGE

DSNU050I 319 18:06:01.49 DSNUGUTC - REORG TABLESPACE DSN00242.XBKRO000
DSNU251I 319 18:06:01.58 DSNURULD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS UNLOADED=1 FOR TABLESPACE
DSN00242.XBKRO000 PART 1
DSNU252I 319 18:06:01.58 DSNURULD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS UNLOADED=1 FOR TABLESPACE
DSN00242.XBKRO000
DSNU250I 319 18:06:01.58 DSNURULD - UNLOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU303I -DBOB 319 18:06:01.97 DSNURWT - (RE)LOAD PHASE STATISTICS - NUMBER OF RECORDS=1 FOR TABLE
XMLR4.XBK_TO_CSTMRT STMT PART=1
DSNU304I -DBOB 319 18:06:01.97 DSNURWT - (RE)LOAD PHASE STATISTICS - NUMBER OF RECORDS=1 FOR TABLE
XMLR4.XBK_TO_CSTMRT STMT
DSNU302I 319 18:06:01.97 DSNURILD - (RE)LOAD PHASE STATISTICS - NUMBER OF INPUT RECORDS PROCESSED=1
DSNU300I 319 18:06:01.97 DSNURILD - (RE)LOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU042I 319 18:06:01.97 DSNUGSOR - SORT PHASE STATISTICS -
      NUMBER OF RECORDS=1
      ELAPSED TIME=00:00:00
DSNU349I -DBOB 319 18:06:02.06 DSNURBXA - BUILD PHASE STATISTICS - NUMBER OF KEYS=1 FOR INDEX
XMLR4.I_NODEIDX BK TO_CSTMRT STMT
DSNU258I 319 18:06:02.06 DSNURBXD - BUILD PHASE STATISTICS - NUMBER OF INDEXES=1
DSNU259I 319 18:06:02.06 DSNURBXD - BUILD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU421I 319 18:06:02.08 DSNUGFUM - START OF DFSMS MESSAGES

1PAGE 0001 5695-DF175 DFSMSDSS V1R12.0 DATA SET SERVICES 2010.319 18:06
- COPY DATASET(INCLUDE( -
  DBOBD.DSNDBC.DSN00242.XBKRO000.I0001.A001 )) -
  RENAMEU( -
    (DBOBD.DSNDBC.DSN00242.XBKRO000.I0001.A001 , -
    DBOBI.DSN00242.XBKRO000.N00001.C2RPDUYM )) -
  REPUNC ALLDATA(*) ALLEXCP CANCELERROR SHARE -
  WRITECHECK TOLERATE(ENQF)
ADR101I (R/I)-RI01 (01), TASKID 001 HAS BEEN ASSIGNED TO COMMAND 'COPY '
ADR109I (R/I)-RI01 (01), 2010.319 18:06:02 INITIAL SCAN OF USER CONTROL STATEMENTS COMPLETED
ADRO50I (001)-PRIME(01), DFSMSDSS INVOKED VIA APPLICATION INTERFACE
ADRO16I (001)-PRIME(01), RACF LOGGING OPTION IN EFFECT FOR THIS TASK
OADR006I (001)-STEND(01), 2010.319 18:06:02 EXECUTION BEGINS
OADR711I (001)-NEWS(01), DATA SET DBOBD.DSNDBC.DSN00242.XBKRO000.I0001.A001 HAS BEEN ALLOCATED WITH NEWNAME
      DBOBI.DSN00242.XBKRO000.N00001.C2RPDUYM USING STORCLAS DBOBDATA, DATACLAS DBOB, AND
MGMTCLAS MCDB22
OADR806I (001)-TOMI (03), DATA SET DBOBD.DSNDBC.DSN00242.XBKRO000.I0001.A001 COPIED USING A FAST REPLICATION
FUNCTION
OADR801I (001)-DDDS (01), DATA SET FILTERING IS COMPLETE. 1 OF 1 DATA SETS WERE SELECTED: 0 FAILED
SERIALIZATION AND 0 FAILED FOR OTHER REASONS
OADR454I (001)-DDDS (01), THE FOLLOWING DATA SETS WERE SUCCESSFULLY PROCESSED
0      DBOBD.DSNDBC.DSN00242.XBKRO000.I0001.A001
OADR006I (001)-STEND(02), 2010.319 18:06:02 EXECUTION ENDS
OADR013I (001)-CLTSK(01), 2010.319 18:06:02 TASK COMPLETED WITH RETURN CODE 0000
OADR012I (SCH)-DSSU (01), 2010.319 18:06:02 DFSMSDSS PROCESSING COMPLETE. HIGHEST RETURN CODE IS 0000
DSNU422I 319 18:06:02.27 DSNUGFCD - END OF DFSMS MESSAGE

DSNU010I 319 18:06:02.29 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0

```

The REORG table space utility reorganizes also the DOCID index (and any user defined indexes) when reorganizing the base table space, and the NODEID index when reorganizing the XML table space.

You can reorganize only the indexes using the REORG INDEX utility.

9.13 REPAIR

You can use the REPAIR utility on XML objects.

You can use the REPAIR utility to:

- ▶ Set the status of an XML column to invalid.
- ▶ Delete a corrupted XML document and its NODEID index entries.

The most common use for the REPAIR utility for XML objects is to take corrective action after you run CHECK DATA with SHRLEVEL CHANGE on a table space with XML columns. CHECK DATA with SHRLEVEL CHANGE operates on shadow data sets, so it does not modify XML columns or XML table spaces. Instead, CHECK DATA generates REPAIR statements that you can run to delete invalid XML documents and to mark the corresponding XML columns as invalid.

Refer to 10.5, “Diagnostics” on page 238 for examples of invoking the REPAIR utility when diagnosing problems with XML data.

9.14 REPORT

When you specify REPORT TABLESPACESET, the output report includes XML objects in the list of members in the table space set.

The sample output in Example 9-20 shows a table space set for a table that contains an XML column:

Example 9-20 REPORT utility JCL (with TABLESPACESET option) and output

```
//XMLR4LD JOB (999,POK),'DB0B',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DB0BM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DB0B,UID='TEMP',UTPROC=' '
//DSNUPROC.SYSIN DD *
REPORT TABLESPACESET DSN00242.BKRTORCS

1DSNU000I 315 16:50:26.26 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 315 16:50:26.28 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 315 16:50:26.28 DSNUGUTC - REPORT TABLESPACESET DSN00242.BKRTORCS
DSNU587I -DB0B 315 16:50:26.29 DSNUPSET - REPORT TABLESPACE SET WITH TABLESPACE
DSN00242.BKRTORCS
```

TABLESPACE SET REPORT:

```
TABLESPACE      : DSN00242.BKRTORCS
TABLE           : XMLR4.BK_TO_CSTMR_STMT
INDEXSPACE     : DSN00242.IRDOCIDB
INDEX          : XMLR4.I_DOCIDBK_TO_CSTMR_STMT
```

XML TABLESPACE SET REPORT:

```
TABLESPACE      : DSN00242.BKRTORCS

BASE TABLE     : XMLR4.BK_TO_CSTMR_STMT
COLUMN          : BK_TO_CSTMR_STMT
```

```
XML TABLESPACE      : DSN00242.XBKR0000
XML TABLE           : XMLR4.XBK_TO_CSTM_ STMT
XML NODEID INDEXSPACE: DSN00242.IRNODEID
XML NODEID INDEX     : XMLR4.I_NODEIDXBK_TO_CSTM_ STMT
```

```
DSNU580I    315 16:50:26.29 DSNUPORT - REPORT UTILITY COMPLETE - ELAPSED TIME=00:00:00
```

NOTE: The same result is produced if you use the following:
REPORT TABLESPACESET DSN00242.XBKR0000

When you specify REPORT RECOVERY, the output report includes recovery related information. Use REPORT RECOVERY to find information that is necessary for recovering a table space, index, or a table space and all of its indexes. This is particularly useful for point-in-time recovery.

Example 9-21 shows the JCL for the REPORT utility with the RECOVERY option for the base table space and the output of the utility run.

Example 9-21 REPORT utility JCL (with RECOVERY option for base table space) and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
REPORT RECOVERY TABLESPACE DSN00242.BKRTORCS

1DSNU000I    315 17:02:40.42 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I    315 17:02:40.45 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I    315 17:02:40.45 DSNUGUTC - REPORT RECOVERY TABLESPACE DSN00242.BKRTORCS
DSNU581I    -DBOB 315 17:02:40.45 DSNUPREC - REPORT RECOVERY TABLESPACE DSN00242.BKRTORCS
DSNU593I    -DBOB 315 17:02:40.46 DSNUPREC - REPORT RECOVERY ENVIRONMENT RECORD:
             MINIMUM RBA: 000000000000
             MAXIMUM RBA: FFFFFFFF0000
             MIGRATING RBA: 000000000000
DSNU582I    -DBOB 315 17:02:40.46 DSNUPPCP - REPORT RECOVERY TABLESPACE DSN00242.BKRTORCS
SYSCOPY ROWS AND SYSTEM LEVEL BACKUPS
TIMESTAMP = 2010-11-04-22.02.43.096801, IC TYPE = *C*, SHR LVL = , DSNUM = 0000,
START LRSN =000066569186
DEV TYPE = , IC BACK = , STYPE = L, FILE SEQ = 0000,
PIT LRSN = 000000000000
LOW DSNUM = 0000, HIGH DSNUM = 0000, OLDEST VERSION = 0000, LOGICAL PART = 0000,
LOGGED = Y, TTYPE =
JOBNAME = , AUTHID = , COPYPAGESF = -1.0E+00
NPAGESF = -1.0E+00 , CPAGESF = -1.0E+00
DSNAME = DSN00242.BKRTORCS , MEMBER NAME = ,
INSTANCE = 01, RELCREATED = M
.....
.....
TIMESTAMP = 2010-11-10-18.12.48.246523, IC TYPE = F , SHR LVL = R, DSNUM = 0000,
START LRSN =000069667C5E
DEV TYPE = 3390 , IC BACK = , STYPE = , FILE SEQ = 0000,
PIT LRSN = 000000000000
LOW DSNUM = 0001, HIGH DSNUM = 0001, OLDEST VERSION = 0000, LOGICAL PART = 0000,
LOGGED = Y, TTYPE =
JOBNAME = XMLR4LD , AUTHID = XMLR4 , COPYPAGESF = 3.0E+00
NPAGESF = 3.4E+01 , CPAGESF = 0.0E0
```



```

DSNAME = DSN00242.BKRTORCS.F.D2010314.T231248.COPY , MEMBER NAME = ,
INSTANCE = 01, RELCREATED = M

TIMESTAMP = 2010-11-10-18.14.25.344136, IC TYPE = I , SHR LVL = R, DSNUM = 0000,
START LRSN =0000696BCB4C
DEV TYPE = 3390 , IC BACK = , STYPE = , FILE SEQ = 0000,
PIT LRSN = 000000000000
LOW DSNUM = 0001, HIGH DSNUM = 0001, OLDEST VERSION = 0000, LOGICAL PART = 0000,
LOGGED = Y, TTYPE =
JOBNAME = XMLR4LD , AUTHID = XMLR4 , COPYPAGESF = 3.0E+00
NPAGESF = 3.4E+01 , CPAGESF = 2.0E+00
DSNAME = DSN00242.BKRTORCS.I.D2010314.T231425.COPY , MEMBER NAME = ,
INSTANCE = 01, RELCREATED = M

TIMESTAMP = 2010-11-10-18.30.01.134384, IC TYPE = F , SHR LVL = R, DSNUM = 0000,
START LRSN =0000696BCB4C
DEV TYPE = 3390 , IC BACK = , STYPE = , FILE SEQ = 0000,
PIT LRSN = 000000000000
LOW DSNUM = 0000, HIGH DSNUM = 0000, OLDEST VERSION = 0000, LOGICAL PART = 0000,
LOGGED = Y, TTYPE =
JOBNAME = XMLR4LD , AUTHID = XMLR4 , COPYPAGESF = 3.0E+00
NPAGESF = 3.4E+01 , CPAGESF = 0.0E0
DSNAME = DSN00242.BKRTORCS.F.D2010314.T233001.COPY , MEMBER NAME = ,
INSTANCE = 01, RELCREATED = M
.....
.....
.....
DSNU580I 315 17:02:40.46 DSNUPORT - REPORT UTILITY COMPLETE - ELAPSED TIME=00:00:00

DSNU010I 315 17:02:40.46 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0

```

Example 9-22 shows the JCL for the REPORT utility with the RECOVERY option for the XML table space and the output of the utility run.

Example 9-22 REPORT utility JCL (with TRECOVERY option for XML table space) and output

```

//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=OM
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
REPORT RECOVERY TABLESPACE DSN00242.XBKRO000

1DSNU000I 315 17:07:09.65 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 315 17:07:09.68 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 315 17:07:09.68 DSNUGUTC - REPORT RECOVERY TABLESPACE DSN00242.XBKRO000
DSNU581I -DBOB 315 17:07:09.68 DSNUPREC - REPORT RECOVERY TABLESPACE DSN00242.XBKRO000
DSNU593I -DBOB 315 17:07:09.69 DSNUPREC - REPORT RECOVERY ENVIRONMENT RECORD:
          MINIMUM RBA: 000000000000
          MAXIMUM RBA: FFFFFFFFFFFF
          MIGRATING RBA: 000000000000
DSNU582I -DBOB 315 17:07:09.69 DSNUPPCP - REPORT RECOVERY TABLESPACE DSN00242.XBKRO000
SYSCOPY ROWS AND SYSTEM LEVEL BACKUPS
TIMESTAMP = 2010-11-04-22.02.43.211430, IC TYPE = *C*, SHR LVL = , DSNUM = 0000,
START LRSN =00006656ED79
DEV TYPE = , IC BACK = , STYPE = L, FILE SEQ = 0000,
PIT LRSN = 000000000000
LOW DSNUM = 0000, HIGH DSNUM = 0000, OLDEST VERSION = 0000, LOGICAL PART = 0000,
LOGGED = Y, TTYPE =

```

```

JOBNAME = , AUTHID = , COPYPAGESF = -1.0E+00
NPAGESF = -1.0E+00 , CPAGESF = -1.0E+00
DSNAME = DSN00242.XBKR0000 , MEMBER NAME = ,
INSTANCE = 01, RELCREATED = M
.....
.....
TIMESTAMP = 2010-11-10-18.12.48.307770, IC TYPE = F , SHR LVL = R, DSNUM = 0000,
START LRSN =000069667C5E
DEV TYPE = 3390 , IC BACK = , STYPE = , FILE SEQ = 0000,
PIT LRSN = 000000000000
LOW DSNUM = 0001, HIGH DSNUM = 0001, OLDEST VERSION = 0000, LOGICAL PART = 0000,
LOGGED = Y, TTYPE =
JOBNAME = XMLR4LD , AUTHID = XMLR4 , COPYPAGESF = 3.0E+00
NPAGESF = 4.5E+01 , CPAGESF = 0.0E0
DSNAME = DSN00242.XBKR0000.F.D2010314.T231248.COPY , MEMBER NAME = ,
INSTANCE = 01, RELCREATED = M

TIMESTAMP = 2010-11-10-18.14.25.421739, IC TYPE = I , SHR LVL = R, DSNUM = 0000,
START LRSN =0000696BCB4C
DEV TYPE = 3390 , IC BACK = , STYPE = , FILE SEQ = 0000,
PIT LRSN = 000000000000
LOW DSNUM = 0001, HIGH DSNUM = 0001, OLDEST VERSION = 0000, LOGICAL PART = 0000,
LOGGED = Y, TTYPE =
JOBNAME = XMLR4LD , AUTHID = XMLR4 , COPYPAGESF = 3.0E+00
NPAGESF = 4.5E+01 , CPAGESF = 2.0E+00
DSNAME = DSN00242.XBKR0000.I.D2010314.T231425.COPY , MEMBER NAME = ,
INSTANCE = 01, RELCREATED = M

TIMESTAMP = 2010-11-10-18.30.01.235727, IC TYPE = F , SHR LVL = R, DSNUM = 0000,
START LRSN =0000696BCB4C
DEV TYPE = 3390 , IC BACK = , STYPE = , FILE SEQ = 0000,
PIT LRSN = 000000000000
LOW DSNUM = 0000, HIGH DSNUM = 0000, OLDEST VERSION = 0000, LOGICAL PART = 0000,
LOGGED = Y, TTYPE =
JOBNAME = XMLR4LD , AUTHID = XMLR4 , COPYPAGESF = 3.0E+00
NPAGESF = 4.5E+01 , CPAGESF = 0.0E0
DSNAME = DSN00242.XBKR0000.F.D2010314.T233001.COPY , MEMBER NAME = ,
INSTANCE = 01, RELCREATED = M
.....
.....
DSNU580I 315 17:07:09.69 DSNUPORT - REPORT UTILITY COMPLETE - ELAPSED TIME=00:00:00
DSNU010I 315 17:07:09.70 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0

```

We do not show the complete output of the REPORT utility, instead show only the relevant entries to demonstrate partial recovery of the table space.

We took a full image copy of the base table space and XML table space (see Example 9-3 on page 188), did a partial update to the XML document (see Figure 9-5 on page 190), took an incremental image copy (see Example 9-4 on page 190), and merged the full and incremental image copies (see Example 9-13 on page 202). The REPORT utility output shows the entries for the full, incremental, and merged full image copy for the base table space in Example 9-21 on page 216 and for the XML table space in Example 9-22 on page 217.

9.15 RUNSTATS

You can use the RUNSTATS utility to gather statistics for XML objects. RUNSTATS TABLESPACE ignores the following keywords for XML table spaces:

- ▶ COLGROUP
- ▶ FREQVAL MOST/LEAST/BOTH
- ▶ HISTOGRAM

RUNSTATS INDEX ignores the following keywords for XML indexes or NODEID indexes:

- ▶ KEYCARD
- ▶ FREQVAL MOST/LEAST/BOTH
- ▶ HISTOGRAM

XML indexes are related to XML tables, and not to the associated base tables. If you specify a base table space and an XML index in the same RUNSTATS control statement, DB2 generates an error. When you run RUNSTATS against a base table, RUNSTATS collects statistics only for indexes on the base table, including the document ID index.

RUNSTATS TABLESPACE does not collect histogram statistics for XML table spaces. RUNSTATS INDEX does not collect histogram statistics for XML NODEID indexes or XML indexes.

Example 9-23 shows the JCL for the RUNSTATS utility for the base table space (and XML table space) and the output of the utility run.

Example 9-23 RUNSTATS utility JCL and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=OM
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSIN DD *
RUNSTATS TABLESPACE DSN00242.BKRTORCS INDEX (ALL)
RUNSTATS TABLESPACE DSN00242.XBKRO000 INDEX (ALL)

1DSNU000I 319 19:06:31.28 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 319 19:06:31.33 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 319 19:06:31.33 DSNUGUTC - RUNSTATS TABLESPACE DSN00242.BKRTORCS INDEX(ALL)
DSNU610I -DBOB 319 19:06:31.40 DSNUSUTP - SYSTABLEPART CATALOG UPDATE FOR
DSN00242.BKRTORCS SUCCESSFUL
DSNU610I -DBOB 319 19:06:31.40 DSNUSUPT - SYSTABSTATS CATALOG UPDATE FOR
XMLR4.BK_TO_CSTMRT_STMT SUCCESSFUL
DSNU610I -DBOB 319 19:06:31.40 DSNUSUTB - SYSTABLES CATALOG UPDATE FOR
XMLR4.BK_TO_CSTMRT_STMT SUCCESSFUL
DSNU610I -DBOB 319 19:06:31.40 DSNUSUTS - SYSTABLESPACE CATALOG UPDATE FOR
DSN00242.BKRTORCS SUCCESSFUL
DSNU610I -DBOB 319 19:06:31.41 DSNUSUIP - SYSINDEXPART CATALOG UPDATE FOR
XMLR4.I_DOCIDBK_TO_CSTMRT_STMT SUCCESSFUL
DSNU610I -DBOB 319 19:06:31.41 DSNUSUCO - SYSCOLUMNS CATALOG UPDATE FOR
XMLR4.I_DOCIDBK_TO_CSTMRT_STMT SUCCESSFUL
DSNU610I -DBOB 319 19:06:31.41 DSNUSUIX - SYSINDEXES CATALOG UPDATE FOR
XMLR4.I_DOCIDBK_TO_CSTMRT_STMT SUCCESSFUL
DSNU610I -DBOB 319 19:06:31.41 DSNUSUCD - SYSCOLDIST CATALOG UPDATE FOR
XMLR4.I_DOCIDBK_TO_CSTMRT_STMT SUCCESSFUL
DSNU620I -DBOB 319 19:06:31.41 DSNUSEOF - RUNSTATS CATALOG TIMESTAMP =
2010-11-15-19.06.31.341899
```

```

ODSNU050I    319 19:06:31.42 DSNUGUTC -  RUNSTATS TABLESPACE DSN00242.XBKRO000 INDEX(ALL)
DSNU610I    -DBOB 319 19:06:31.48 DSNUSUTP - SYSTABLEPART CATALOG UPDATE FOR
DSN00242.XBKRO000 SUCCESSFUL
DSNU610I    -DBOB 319 19:06:31.48 DSNUSUPT - SYSTABSTATS CATALOG UPDATE FOR
XMLR4.XBK_TO_CSTMR_STMT SUCCESSFUL
DSNU610I    -DBOB 319 19:06:31.48 DSNUSUTB - SYSTABLES CATALOG UPDATE FOR
XMLR4.XBK_TO_CSTMR_STMT SUCCESSFUL
DSNU610I    -DBOB 319 19:06:31.48 DSNUSUTS - SYSTABLESPACE CATALOG UPDATE FOR
DSN00242.XBKRO000 SUCCESSFUL
DSNU610I    -DBOB 319 19:06:31.49 DSNUSUIP - SYSINDEXPART CATALOG UPDATE FOR
XMLR4.I_NODEIDX BK_TO_CSTMR_STMT SUCCESSFUL
DSNU610I    -DBOB 319 19:06:31.49 DSNUSUKT - SYSKEYTARGETS CATALOG UPDATE FOR
XMLR4.I_NODEIDX BK_TO_CSTMR_STMT SUCCESSFUL
DSNU610I    -DBOB 319 19:06:31.49 DSNUSUIX - SYSINDEXES CATALOG UPDATE FOR
XMLR4.I_NODEIDX BK_TO_CSTMR_STMT SUCCESSFUL
DSNU620I    -DBOB 319 19:06:31.49 DSNUSEOF - RUNSTATS CATALOG TIMESTAMP =
2010-11-15-19.06.31.425894

DSNU010I    319 19:06:31.50 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0

```

9.16 UNLOAD

You can unload XML data to output data records or to separate files.

XML columns can be unloaded with either of the following methods:

- ▶ The XML column can be unloaded to the output records. XML column value can be placed in the OUTPUT record with or without any other unloading column values. The output record can be in delimited or non-delimited format. For a non-delimited format, the XML column is handled like a variable character field with a 2-byte length preceding the XML value. For a delimited format there are no length bytes present. If the total output record length is more than 32 KB, unload the record in spanned record format by specifying the SPANNED YES option.
- ▶ The XML column can be unloaded to a separate file whether the XML column length is less than 32 KB or not.

The output data can be in the textual XML format or the binary XML format. Data that is unloaded can be in the delimited or non-delimited format.

To unload XML data directly to output record specify XML as the output field type. If the output is a non-delimited format, a 2-byte length will precede the value of the XML. For delimited output, no length field is present. XML is the only acceptable field type when unloading the XML directly to the output record. No data type conversion applies and you cannot specify FROMCOPY.

If the input data is in Extensible Dynamic Binary XML DB2 Client/Server Binary XML Format (binary XML format), you need to specify BINARYXML. To unload XML data to a separate file:

- ▶ In the UNLOAD utility control statement, specify BLOBF, CLOBF or DBCLOBF. These keywords indicate that the output column contains the name of a file to which the XML value is to be unloaded. Also specify either CHAR or VARCHAR instead of XML. Do not specify FROMCOPY.
- ▶ Use the template control statement to create the XML output file and filename. If data sets are not created and the DSN type is not specified on the template, UNLOAD will use PDS

as the data set type. PDS has a limit of single volume. The output file uses multiple volumes, so you must specify HFS as the DSN type.

In the UNLOAD statement, specify the base table space. You cannot specify the XML table space.

Example 9-24 shows the JCL for the UNLOAD utility and the output of the utility run.

Example 9-24 UNLOAD utility JCL and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSREC DD DSN=XMLR4.UNLOAD
// DISP=(MOD,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSPUNCH DD DSN=XMLR4.LOADCTL
// DISP=(MOD,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSIN DD *
UNLOAD TABLESPACE DSN00242.BKRTORCS
FROM TABLE XMLR4.BK_TO_CSTMRTMT
(BK_TO_CSTMRTMT POSITION(*) XML)

1DSNU000I 301 18:01:07.19 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 301 18:01:07.21 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 301 18:01:07.22 DSNUGUTC - UNLOAD TABLESPACE DSN00242.BKRTORCS
DSNU650I -DBOB 301 18:01:07.22 DSNUUGMS - FROM TABLE XMLR4.BK_TO_CSTMRTMT
DSNU650I -DBOB 301 18:01:07.22 DSNUUGMS - (BK_TO_CSTMRTMT POSITION(*) XML)
DSNU253I 301 18:01:07.24 DSNUUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
UNLOADED=1 FOR TABLE XMLR4.BK_TO_CSTMRTMT
DSNU252I 301 18:01:07.24 DSNUUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
UNLOADED=1 FOR TABLESPACE DSN00242.BKRTORCS
DSNU250I 301 18:01:07.24 DSNUUNLD - UNLOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU010I 301 18:01:07.25 DSNUGBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
```

NOTE: SYSREC shows the name of the data set which receives the XML document unloaded by UNLOAD utility.

SYSPUNCH shows the name of the data set which has the LOAD utility control statements generated by UNLOAD utility.

The first few characters of the XML document in SYSREC data set are shown below:

```
7<?xml version="1.0" encoding="IBM037"?><Document xmlns:xsi="http://
0001F46A994A89A89977F4F748989889877CCDFFF7664C98A989A4A999A7AA8778AA9766
03017CF74305592965EF1B0F055364957EF924037FFEC46344553074352A729EF8337A11
A 2-byte length precedes the XML value.
```

The LOAD utility control statement in SYSPUNCH data set is shown below:

```
LOAD DATA INDDN SYSREC LOG NO RESUME YES
EBCDIC CCSID(00037,00000,00000)
INTO TABLE "XMLR4"."BK_TO_CSTMRTMT"
WHEN(00001:00002) = X'0003'
NUMRECS 1
( "BK_TO_CSTMRTMT"
POSITION( *) XML PRESERVE WHITESPACE
)
```

UNLOAD utility using file reference variable

Example 9-25 shows the JCL for the UNLOAD utility using file reference variable and the output of the utility run.

Example 9-25 UNLOAD utility JCL (using file reference variable) and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=OM
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSREC DD DSN=XMLR4.UNLOAD1,
// DISP=(MOD,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSPUNCH DD DSN=XMLR4.LOADCTL1,
// DISP=(MOD,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSIN DD *
TEMPLATE TCLOBF UNIT(SYSDA) DISP(MOD,CATLG,DELETE)
DSN(&USERID..&DB..&TS..&TI..UFILEREf)
UNLOAD TABLESPACE DSN00242.BKRTORCS
FROM TABLE XMLR4.BK_TO_CSTMRTMT
(BK_TO_CSTMRTMT POSITION(*) VARCHAR CLOBF TCLOBF)

1DSNU000I 301 19:12:35.80 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 301 19:12:35.82 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
ODSNU050I 301 19:12:35.83 DSNUGUTC - TEMPLATE TCLOBF UNIT(SYSDA) DISP(MOD, CATLG,
DELETE) DSN(&USERID..&DB..&TS..&TI..UFILEREf)
DSNU1035I 301 19:12:35.83 DSNUJTDR - TEMPLATE STATEMENT PROCESSED SUCCESSFULLY
ODSNU050I 301 19:12:35.83 DSNUGUTC - UNLOAD TABLESPACE DSN00242.BKRTORCS
DSNU650I -DBOB 301 19:12:35.83 DSNUUGMS - FROM TABLE XMLR4.BK_TO_CSTMRTMT
DSNU650I -DBOB 301 19:12:35.83 DSNUUGMS - (BK_TO_CSTMRTMT POSITION(*) VARCHAR CLOBF
TCLOBF)
DSNU1038I 301 19:12:35.89 DSNUGDYN - DATASET ALLOCATED. TEMPLATE=TCLOBF
DDNAME=SYS0001
DSN=XMLR4.DSN00242.XBKRO000.T231235.UFILEREf
DSNU253I 301 19:12:35.92 DSNUUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
UNLOADED=1 FOR TABLE XMLR4.BK_TO_CSTMRTMT
DSNU252I 301 19:12:35.92 DSNUUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
UNLOADED=1 FOR TABLESPACE DSN00242.BKRTORCS
DSNU250I 301 19:12:35.93 DSNUUNLD - UNLOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU010I 301 19:12:35.94 DSNUBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
```

NOTE: SYSREC shows the name of the data set derived from the Template definition which has the XML document as shown below:

```
XMLR4.DSN00242.XBKRO000.T231235.UFILEREf(C1XSJ2ZY)
```

The first few characters of the XML document are shown below:

```
<?xml version="1.0" encoding="IBM037"?><Document xmlns:xsi="http://www
46A994A89A89977F4F748989889877CCDFFF7664C98A989A4A999A7AA8778AA9766AAA
CF74305592965EF1B0F055364957EF924037FFEC46344553074352A729EF8337A1166A
```

The XML document starts from position 1.

The LOAD utility control statement in SYSPUNCH data set is shown below:

```
LOAD DATA INDDN SYSREC LOG NO RESUME YES
```

```
EBCDIC CCSID(00037,00000,00000)
```

```
INTO TABLE
```

```
"XMLR4"."BK_TO_CSTMRTMT"
```

```
WHEN(00001:00002) = X'0003'
```

```
NUMRECS 1
```

```
( "BK_TO_CSTMRTMT"
  POSITION( 00003:00259) VARCHAR CLOBF          PRESERVE WHITESPACE
)
```

UNLOAD utility to unload XML data in binary

Example 9-26 shows the JCL for the UNLOAD utility to unload XML data in binary and the output of the utility run.

Example 9-26 UNLOAD utility JCL (to unload XML data in binary) and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=OM
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSREC DD DSN=XMLR4.UNLOAD.XML.BINARY.RECOO,
// DISP=(MOD,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSPUNCH DD DSN=XMLR4.UNLOAD.XML.BINARY.PUNCH,
// DISP=(MOD,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSIN DD *
UNLOAD DATA
FROM TABLE XMLR4.BK_TO_CSTMRTMT
 (BK_TO_CSTMRTMT XML BINARYXML)

1DSNU000I 308 14:17:16.66 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 308 14:17:16.69 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
0DSNU050I 308 14:17:16.69 DSNUGUTC - UNLOAD DATA
DSNU650I -DBOB 308 14:17:16.69 DSNUUGMS - FROM TABLE XMLR4.BK_TO_CSTMRTMT
DSNU650I -DBOB 308 14:17:16.69 DSNUUGMS - (BK_TO_CSTMRTMT XML BINARYXML)
DSNU253I 308 14:17:16.76 DSNUUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
UNLOADED=1 FOR TABLE XMLR4.BK_TO_CSTMRTMT
DSNU252I 308 14:17:16.76 DSNUUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
UNLOADED=1 FOR TABLESPACE DSN00242.BKRTORCS
DSNU250I 308 14:17:16.77 DSNUUNLD - UNLOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
```

NOTE: SYSREC shows the name of the data set which receives the XML document unloaded by UNLOAD utility.

SYSPUNCH shows the name of the data set which has the LOAD utility control statements generated by UNLOAD utility.

The first few characters of the XML document in SYSREC data set are shown below:

```
....-.....ñ.İËñ.ñ.ÇËËø...İİ.İ..?ËÅ.....İ (<ëÄÇÄ_/ .Ñ>ËË/>ÄÄÄñ.İË>.ÑË?.ËËÄ.ÑË?.
0003C3000000407761426777322777273267623332544566666626677666666427763676377636763
03D4AB510002938398998440AFF777E73EF27F2001F8DC3385D1D9E341E3569E52EA93FA344A93FA
A 2-byte length precedes the XML value.
```

The LOAD utility control statement in SYSPUNCH data set is shown below:

```
LOAD DATA INDDN SYSREC LOG NO RESUME YES
EBCDIC CCSID(00037,00000,00000)
INTO TABLE "XMLR4"."BK_TO_CSTMRTMT"
WHEN(00001:00002) = X'0003'
NUMRECS 1
("BK_TO_CSTMRTMT"
 POSITION( *) XML PRESERVE WHITESPACE BINARYXML)
```

UNLOAD utility to unload into a VBS data set in spanned record format

If you want to unload data from a table that has large LOB or XML fields, consider unloading the data in spanned record format to improve performance of read-write operations.

When you unload data in spanned record format, all LOB and XML data for a given table space or table space partition can be written to an individual sequential file. This file can reside on DASD and can span multiple volumes. Having such a single sequential file can improve the performance of read-write operations.

To unload data in spanned record format, specify the SPANNED YES option. Specify in the field specification list that all LOB and XML data are to be at the end of the record.

Example 9-27 shows the JCL for the UNLOAD utility to unload XML data in spanned record format and the output of the utility run.

Example 9-27 UNLOAD utility JCL (to unload XML data in spanned record format) and output

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=XMLR4,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//UTIL EXEC DSNUPROC,SYSTEM=DBOB,UID='TEMP',UTPROC=''
//DSNUPROC.SYSREC DD DSN=XMLR4.DSN00242.XBKRO000.T214620.UFILEREF.NEW,
// DISP=(MOD,CATLG)
//DSNUPROC.SYSPUNCH DD DSN=XMLR4.UNLOAD.SCENARIO.NEW.PUNCH3,
// DISP=(MOD,CATLG),
// SPACE=(16384,(20,20),,,ROUND),
// UNIT=SYSDA
//DSNUPROC.SYSIN DD *
UNLOAD DATA SPANNED YES FROM TABLE XMLR4.BK_TO_CSTMR_STMT
      (MSG_ID VARCHAR,
      MSG_CRE_DT_TM TIMESTAMP WITH TIME ZONE EXTERNAL,
      BK_TO_CSTMR_STMT XML)
```

NOTE: The record format of the SYSREC data set is **VBS**.

```
1DSNU000I 309 18:59:48.56 DSNUGUTC - OUTPUT START FOR UTILITY, UTILID = TEMP
DSNU1044I 309 18:59:48.58 DSNUGTIS - PROCESSING SYSIN AS EBCDIC
0DSNU050I 309 18:59:48.59 DSNUGUTC - UNLOAD DATA SPANNED YES
DSNU650I -DBOB 309 18:59:48.59 DSNUUGMS - FROM TABLE XMLR4.BK_TO_CSTMR_STMT
DSNU650I -DBOB 309 18:59:48.59 DSNUUGMS - (MSG_ID VARCHAR,
DSNU650I -DBOB 309 18:59:48.59 DSNUUGMS - MSG_CRE_DT_TM TIMESTAMP WITH TIME ZONE
EXTERNAL,
DSNU650I -DBOB 309 18:59:48.59 DSNUUGMS - BK_TO_CSTMR_STMT XML)
DSNU253I 309 18:59:48.65 DSNUUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
UNLOADED=1 FOR TABLE XMLR4.BK_TO_CSTMR_STMT
DSNU252I 309 18:59:48.65 DSNUUNLD - UNLOAD PHASE STATISTICS - NUMBER OF RECORDS
UNLOADED=1 FOR TABLESPACE DSN00242.BKRTORCS
DSNU250I 309 18:59:48.66 DSNUUNLD - UNLOAD PHASE COMPLETE, ELAPSED TIME=00:00:00
DSNU010I 309 18:59:48.66 DSNUBAC - UTILITY EXECUTION COMPLETE, HIGHEST RETURN CODE=0
```

From performance point of view, in general HFS is much better than PDS for file references. Binary XML is significantly better than textual XML. The best performing combination is SPANNED with BINARYXML.

9.17 DSNTIAUL

Like the UNLOAD utility, the DSNTIAUL sample program also provides two ways to handle XML data:

- ▶ Unload XML columns as normal data columns to the SYSRECxx file.
- ▶ Unload XML columns to a separate file.

Unload LOB data as normal data columns (SQL parameter)

As the maximum record length of a sequential file in z/OS is 32 KB, this method can only be used if the total record size of the data to be unloaded does not exceed 32 KB. The XML fields are unloaded together with the other selected data fields to the output file with a maximum record length of 32 KB.

In most cases, this method is only used if the XML documents in the table are small.

We demonstrate this with the table XMLR4.BK_TO_CSTMR_STMT defined in Example 4-1 on page 52.

In the first case, we execute the JCL as shown in Example 9-28.

Example 9-28 DSNTIAUL with SQL parameter

```
//XMLR4LD JOB (999,POK),'DBOB',CLASS=A,
// MSGCLASS=T,NOTIFY=&SYSUID,TIME=1440,REGION=0M
/*JOBPARM SYSAFF=SC63,L=9999
// JCLLIB ORDER=(DBOBM.PROCLIB)
//JOBLIB DD DSN=DBOBT.SDSNLOAD,DISP=SHR
//DSNTIAUL EXEC PGM=IKJEFT01,DYNAMNBR=20,COND=(4,LT)
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
DSN SYSTEM(DBOB)
RUN PROGRAM(DSNTIAUL) PLAN(DSNTIB10) PARM('SQL') -
LIB('DBOBM.RUNLIB.LOAD')
//SYSPRINT DD SYSOUT=*
//SYSRECOO DD DSN=XMLR4.DBOB.DSN8UNLD.SQL.SYSRECOO,
// DISP=(MOD,CATLG),UNIT=3390,
// SPACE=(CYL,(100,100))
//SYSPUNCH DD DSN=XMLR4.DBOB.DSN8UNLD.SQL.SYSPUNCH,
// DISP=(MOD,CATLG),UNIT=3390,
// SPACE=(CYL,(1,1))
//SYSIN DD *
SELECT BK_TO_CSTMR_STMT
FROM XMLR4.BK_TO_CSTMR_STMT ;
```

As a result, we get an FB sequential data set SYSRECOO with a LRECL and BLKSIZE of 32,753 bytes with BK_TO_CSTMR_STMT beginning in position 1.

```
DSNT490I SAMPLE DATA UNLOAD PROGRAM
DSNT505I DSNTIAUL OPTIONS USED: SQL
DSNT503I UNLOAD DATA SET SYSPUNCH RECORD LENGTH SET TO 80
DSNT504I UNLOAD DATA SET SYSPUNCH BLOCK SIZE SET TO 27920
DSNT506I INPUT STATEMENT WAS NOT A FULL SELECT ON A SINGLE TABLE. LOAD STATEMENT WILL NEED
MODIFICATION.
DSNT503I UNLOAD DATA SET SYSRECOO RECORD LENGTH SET TO 32760
DSNT504I UNLOAD DATA SET SYSRECOO BLOCK SIZE SET TO 32760
DSNT495I SUCCESSFUL UNLOAD 1 ROWS OF TABLE TBLNAME
The job ends with a return code of RC=4.
```

NOTE:

The first few characters in the unloaded SYSRECOO data set are shown below:

```

...Û<?xml version="1.0" encoding="IBM037"?><Document xmlns:xsi="http://www.w3.or
001F46A994A89A89977F4F748989889877CCDFFF7664C98A989A4A999A7AA8778AA9766AAA4AF499
001CCF74305592965EF1B0F055364957EF924037FFEC46344553074352A729EF8337A11666B63B69
DSNTIAUL generates in the SYSPUNCH data set the following LOAD control statement:
LOAD DATA LOG NO INDDN SYSRECOO INTO TABLE TBLNAME
(BK_TO_CSTMR_STMT      POSITION( 1 ) CLOB )

```

In the SYSIN DD * if you specify all the column names explicitly as shown below:

```

SELECT DB2_GENERATED_DOCID_FOR_XML,
       MSG_ID,
       MSG_CRE_DT_TM,
       BK_TO_CSTMR_STMT
FROM XMLR4.BK_TO_CSTMR_STMT ;
DSNTIAUL issues the messages as above including DSNT506I and generates in the SYSPUNCH data
set the following LOAD control statement:
LOAD DATA LOG NO INDDN SYSRECOO INTO TABLE TBLNAME
(DB2_GENERATED_DOCID_FOR_XML POSITION( 1 ) BIGINT NULLIF( 9)='?',
 MSG_ID      POSITION( 10 ) VARCHAR NULLIF(47)='?',
 MSG_CRE_DT_TM POSITION( 48 ) TIMESTAMP WITH TIME ZONE EXTERNAL(32)
 NULLIF(80)='?',
 BK_TO_CSTMR_STMT      POSITION( 81 ) CLOB )

```

You should edit the LOAD control statement by replacing TBLNAME by the actual table name, CLOB by XML and POSITION (1) by POSITION (3) or POSITION (81) by POSITION (83) depending on how you specify the SELECT statement in DSNTIAUL, include either RESUME YES or REPLACE, and change INDDN SYSRECOO to INDDN SYSREC.

In the SYSIN DD * if you specify SELECT * FROM XMLR4.BK_TO_CSTMR_STMT ; DSNTIAUL issues the messages as above **excluding** DSNT506I, the job ends with a return code of **RC=0**, and generates in the SYSPUNCH data set the following LOAD control statement:

```

LOAD DATA LOG NO INDDN SYSRECOO INTO TABLE BK_TO_CSTMR_STMT
(MSG_ID      POSITION( 1 ) VARCHAR NULLIF(38)='?',
 MSG_CRE_DT_TM POSITION( 39 ) TIMESTAMP WITH TIME ZONE EXTERNAL(32)
 NULLIF(71)='?',
 BK_TO_CSTMR_STMT      POSITION( 72 ) CLOB )

```

Notice the DB2_GENERATED_DOCID_FOR_XML column is not included.

You should edit the LOAD control statement by replacing CLOB by XML and POSITION (72) by POSITION (74), include either RESUME YES or REPLACE, and change INDDN SYSRECOO to INDDN SYSREC.

Unload XML data to a separate file (LOBFIL parameter)

This is the recommended method introduced since DB2 9. With this method, the XML values are unloaded to a different file than the normal SYSRECxx unload files. DSNTIAUL dynamically creates a sequential data set for each XML document to be unloaded. The name of the separate file is stored in the normal SYSRECxx unload files together with the other normal data fields.

DSNTIAUL does this by using the new XML file reference variables introduced since DB2 9. Each XML file has a name of the form <prefix>.Q<i>.C<j>.R<k>, where:

- ▶ <prefix> is a user-specified data set name prefix. <prefix> must conform to the rules for a z/OS physical sequential data set name and cannot exceed 17 characters.
- ▶ Q<i> is the (<i>-1)th query processed by the current DSNTIAUL session. <i> ranges from 0,000,000 to 0,000,099, which corresponds to the limit on the number of queries that can be processed by a single DSNTIAUL session.

You should edit the LOAD control statement to include either RESUME YES or REPLACE, and change INDDN SYSREC00 to INDDN SYSREC.

You can also specify SQL as first PARM:

```
RUN PROGRAM(DSNTIAUL) PLAN(DSNTIB91) PARMS('SQL,LOBFILE(XMLR4)') -
```

And you can use an SQL statement as input to SYSIN:

```
SELECT BK_TO_CSTMRT_STMT FROM XMLR4.BK_TO_CSTMRT_STMT ;
```

In both cases, we get a sequential data set SYSREC00 containing the normal data fields and the names of the XML output files. All of the XML output files have a name of the form XMLR4.Q0000000.C0000006.R000xxxx with xxxx from 0000 to 5882. They are dynamically allocated as sequential files with RECFM=VB,LRECL=27994,BLKSIZE=27998, which is the optimal BLKSIZE for 3390 devices.

9.18 DSN1COPY

You can use DSN1COPY to copy tables from one subsystem to another. When you copy tables from one subsystem to another, you must ensure that the version information on the target subsystem matches the version information on the source subsystem.

Restriction: DB2 XML data is condensed by substituting strings by unique IDs. These unique IDs are stored in the catalog table SYSIBM.SYSXMLSTRINGS and they are not available in the related XML table space. Therefore it is not recommended to copy XML table spaces from one subsystem to another using DSN1COPY.



XML-related tasks for the DBA

The intent of this chapter is to provide the database administrators (DBA) with an overview of what they need to do in connection with XML.

We describe the typical DBA tasks affected by XML by grouping them in the following sections:

- ▶ Tasks regarding system setup
- ▶ Tasks regarding object creation
- ▶ Housekeeping
- ▶ Backup and recovery
- ▶ Diagnostics

Many of the DBA tasks are performed using utilities, and for these we include only a brief overview. Utilities that have additional considerations with the advent of pureXML are covered in Chapter 9, “Utilities with XML” on page 181.

10.1 Tasks regarding system setup

Before creating and using XML objects, you might need to do some configuration and setup, such as set up XML schema repository (XSR) for XML validation, size and allocate a dedicated XML buffer pool and size and adjust the amount of memory available for XML processing.

10.1.1 Setting up the XSR

Before you can do XML schema validation of your XML documents, you need to set up the XSR. This involves creating a set of DB2 tables and indexes that store XML schema information, and a set of stored procedures that operate on the XML schemas that are stored in the tables.

The major steps to set up the XSR are:

1. Define the XSR tables and indexes.

Installation job DSNTIJRT invokes a program that executes the CREATE DATABASE, CREATE TABLESPACE, CREATE TABLE and CREATE INDEX statements for the XML schema repository tables and indexes. After the installation process customizes job DSNTIJRT, you can run DSNTIJRT without further modification to create those tables and indexes. Important: Do not drop these objects after you begin to do XML schema validation. Doing so can cause unexpected behavior.

2. Define the WLM environment and startup procedure for the C language XSR stored procedures.
3. Define the WLM environment and startup procedure for the Java language XSR stored procedure.

Note: The Java stored procedure XSR_COMPLETE needs to run non-APF authorized. This can be accomplished by adding a non-APF authorized data set to the steplib concatenation of the WLM procedure.

However, be aware that other Java stored procedures may need to run APF authorized, so you may need to create a special WLM procedure for XSR_COMPLETE.

4. Define the XML schema repository stored procedures to DB2.
5. Bind the packages for the XML schema repository stored procedures.

Installation job DSNTIJRT invokes a program that binds the packages for the XML schema repository stored procedures. After the installation process customizes job DSNTIJRT, you can run DSNTIJRT without further modification to bind the packages.

6. Bind the packages for the IBM Data Server Driver for JDBC and SQLJ.
7. Test the XML schema repository setup.

For detailed information, reference “Setting up the XML schema repository” in *DB2 10 for z/OS pureXML Guide*, SC19-2981.

10.1.2 Buffer pool for XML

When you create a table with an XML column or alter a table to add an XML column, DB2 creates the XML table space and indexes implicitly. The buffer pool used for XML table spaces is always 16 KB.

The DEFAULT BUFFER POOL FOR USER XML DATA field (TBSBPXML subsystem parameter) specifies the default buffer pool that is to be used for XML table spaces. The default is BP16K0.

You can alter the BUFFERPOOL property of the XML table space, which supports the altering to other 16 KB buffer pools only.

10.1.3 Sizing XMLVALA and XMLVALS

You can use the XMLVALA and XMLVALS subsystem parameters to limit the amount of DB2 virtual storage that is used for XML processing. Because XML values are not fixed in length, and could be very large, DB2 cannot estimate the amount of memory that it needs for processing SQL/XML and XPath queries before run time. DB2 allocates virtual storage at run time based on the size of the XML data. For large XML data, the amount of virtual storage that DB2 requires can grow very large.

If your DB2 subsystem encounters storage constraints because XML values are using too much memory, set the XMLVALA and XMLVALS subsystem parameters:

- ▶ To specify the maximum amount of memory, in KB, for storing XML values for each user, set XMLVALA. The default is 204800 KB.
- ▶ To specify the maximum amount of memory, in MB, for storing XML values for the entire subsystem, set XMLVALS. The default is 10240 MB.

10.1.4 Be up to date with maintenance

In DB2 9, many enhancements have been introduced through APARs, such as XML index for joining and some XMLTABLE performance enhancement. It is recommended to apply these APARs and keep your DB2 system up to date, so you can benefit from these and future enhancements.

As DB2 utilizes XML System Services for parsing and validating XML documents, attention should also be paid to any maintenance offered in this area.

See information APAR II14426 for XML service.

10.2 Tasks regarding object creation

When you create tables with XML columns or add XML columns to existing tables, all the XML objects are created implicitly without any action from the DBA.

However, you need to be aware of the choices that are made for these objects. Some properties may be inherited directly from the base table, some may be inferred from properties of the base table, and some may be dependant on default values or values of DSNZPARMs.

It is a good idea to be aware of how these properties are derived, so you can plan them or alter them after the objects have been created.

We go through the most important properties in this section.

10.2.1 Creation of table with XML columns

You can create a table with XML columns, or alter a table to add one or more XML column. When a table with an XML column is created, an XML table space, document ID index and NODEID index are implicitly created.

For more information about creation of tables with XML columns and the storage structure for XML data, please reference Chapter 4, “Creating and adding XML data” on page 51.

10.2.2 Alteration of implicitly created XML objects

After creating or adding an XML column, you are able to alter implicitly created XML objects. However, you can change only some of the properties for the XML objects. The properties can be altered are listed in Table 10-1.

Table 10-1 Properties can be altered for XML objects

Objects	Description
XML table space	<p>You can alter the following properties:</p> <ul style="list-style-type: none"> ▶ BUFFERPOOL (16 KB buffer pools only) ▶ COMPRESS ▶ PRIQTY ▶ SECQTY ▶ MAXROWS ▶ FREEPAGE ▶ PCTFREE ▶ GBPCACHE ▶ USING STOGROUP ▶ ERASE ▶ LOCKSIZE (The only possible values are XML and TABLESPACE.) ▶ SEGSIZE ▶ DSSIZE ▶ MAXPARTITIONS <p>XML table space attributes that are inherited from the base table space, such as LOG, are implicitly altered if the base table space is altered.</p>
XML table	The ALTER TABLE ALTER PARTITION statement is not supported if the table contains an XML column.
Index	<p>You cannot alter the following properties:</p> <ul style="list-style-type: none"> ▶ CLUSTER ▶ PADDED ▶ ADD COLUMN

10.2.3 Sizing table spaces

Special consideration has to be paid to sizing table spaces when dealing with XML data in range-partitioned table spaces. Recall that the type of the base table space dictates the type of the implicitly created XML table space. The correspondence is shown in Table 10-2.

Table 10-2 Table space types for base and XML tables

Base table space	XML table space
Simple	Partition-by-growth
Segmented	Partition-by-growth
Partitioned	Range-partitioned
Partition-by-growth	Partition-by-growth
Range-partitioned	Range-partitioned

For partition-by-growth XML table spaces, there is no correspondence between the partition in which a particular XML document resides, and the partition of the base table row. The table spaces grow as needed, and independently of each other.

For partitioned and range-partitioned table spaces, the XML document and the base row must reside in corresponding partitions. If the base table rows moves partition, so does the XML document. Therefore, the number of rows fitting into a relational partition is limited by the number of rows that fit into the XML partition.

The DSSIZE of the XML table space is dictated by a combination of the DSSIZE and page size of the base table. The exact values are shown in Table 10-3.

Table 10-3 DSSIZE of the XML table space

DSSIZE of base table space	Page size 4K	Page size 8K	Page size 16K	Page size 32K
1 - 4 GB	4 GB	4 GB	4 GB	4 GB
8 GB	32 GB	16 GB	16 GB	16 GB
16 GB	64 GB	32 GB	16 GB	16 GB
32 GB	64 GB	64 GB	32 GB	16 GB
64 GB	64 GB	64 GB	64 GB	64 GB

To understand what impact the decisions made with respect to size and range of partitions, let us assume that we want to implement the BK_TO_CSTMR_STMT table as a range-partitioned table with one partition per year. The DDL to create this table is shown in Example 10-1.

Example 10-1 Creating a range-partitioned table

```

CREATE TABLESPACE BKSTTS01
  IN BKSTDB01
  USING STOGROUP SYSDEFLT
  DSSIZE 4G
  BUFFERPOOL BPO
  Numparts 3#
CREATE TABLE XMLR2.BK_TO_CSTMR_STMT
  (MSG_ID          VARCHAR(35) FOR SBCS DATA
   WITH DEFAULT NULL,
  MSG_CRE_DT_TM   TIMESTAMP (6)
   WITH DEFAULT NULL,
  BK_TO_CSTMR_STMT XML
   (XMLSCHEMA ID SYSXSR.CAMT_053_001_02)
   NOT NULL)

```

```

IN BKSTDB01.BKSTTS01
PARTITION BY RANGE (MSG_CRE_DT_TM)
(PARTITION 1 ENDING AT ('2009-12-31T23:59:59.99999'),
PARTITION 2 ENDING AT ('2010-12-31T23:59:59.99999'),
PARTITION 3 ENDING AT ('2011-12-31T23:59:59.99999')) #

```

Now assume that the average size of a BankToCustomerStatement is 4K. The base table is very small with a maximum row length of 82 bytes.

If we need to store 10,000 bank statements each year, this yields in round numbers

- ▶ 10,000 * 82 bytes = 800 KB for the base table space
- ▶ 10,000 * 4 KB = 40 MB for the XML table space

In other words, if we just stick with the (default) DSSIZE of 4GB and page size of 4K, we will have sufficient space for the data.

However, what if instead of 10,000 bank statements, we have 2 million each year? In this case, the ballpark estimate becomes

- ▶ 2,000,000 * 82 bytes = 160 MB for the base table space
- ▶ 2,000,000 * 4 KB = 7,8 GB for the XML table space

So even if there is plenty of space for the base table row in each partition using DSSIZE 4G, the limit of the XML table space partition is reached long before the 2 million rows and this imposes a limit on the base table as well. An attempt to insert a row when the XML table space partition is full will result in an SQL code -904, indicating that the XML table space is unavailable.

The solution is either to use a more granular partitioning key, or to choose a combination of DSSIZE and page size of the base table space that will give us a larger DSSIZE for the XML table space.

10.2.4 Compression

Using compression can significantly reduce the amount of disk space needed to store XML data.

To compress data:

1. Specify COMPRESS YES in your XML table space
 - The COMPRESS property of implicitly created XML table space is inherited from base table.
 - To check the COMPRESS property of your table space, query the COMPRESS column of SYSIBM.SYSTABLEPART catalog table
 - You can ALTER TABLESPACE with COMPRESS clause to change the COMPRESS property

2. Run REORG utility

If the XML table space has the COMPRESS YES attribute, the XML data will be compressed

With DB2 10 NFM, you can turn on compression with ALTER any time and the compression dictionary is built when you execute:

- ▶ INSERT statements

- ▶ MERGE statements
- ▶ LOAD SHRLEVEL CHANGE and SHRLEVEL NONE

Additionally, when you INSERT/MERGE/LOAD XML data, a dictionary could be built specifically for the XML table space if the amount of XML data is large enough, and then the new inserted XML data will be compressed.

Note: Real time statistics (RTS) keeps track of the amount of data for the threshold of online compression. Apply the PTF for APAR PM22081 to correct RTS errors when LOAD RESUME is executed on partitioned table spaces.

The compression dictionary is built through these if:

- ▶ The table space or partition is defined with COMPRESS YES
- ▶ The table space or partition has no compression dictionary built yet
- ▶ The amount of data in the table space is large enough to build the compression dictionary.

You can also compress the XML indexes as shown in Example 10-2.

Example 10-2 Creating an XML index with compression

```
CREATE INDEX LEANXML_IX3
ON BK_TO_CSTMRT_STMT(BK_TO_CSTMRT_STMT)
GENERATE KEY USING XMLPATTERN
'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02"#
/Document/BkToCstmrtStmt/GrpHdr/MsgId'
AS SQL VARCHAR(35)
COMPRESS YES#
```

10.2.5 Registration of schemas

An XML schema consists of a set of XML schema documents. You need to register the XML schema to DB2 XSR before using it. You can register an XML schema in any of the following ways:

- ▶ Call the following DB2-supplied stored procedures from a DB2 application program
 - SYSPROC.XSR_REGISTER: Begins registration of an XML schema.
 - SYSPROC.XSR_ADDSCHEMADOC: Adds additional XML schema documents to an XML schema
 - SYSPROC.XSR_COMPLETE: Completes the registration of an XML schema.
- ▶ Invoke the following JDBC method from a Java application program
 - com.ibm.db2.jcc.DB2Connection.registerDB2XmlSchema
- ▶ Invoke the following commands from the Command Line Processor
 - -REGISTER XMLSCHEMA
 - -ADD XMLSCHEMA DOCUMENT
 - -COMPLETE XMLSCHEMA

We show how to register XML schema using CLP in Example 6-2 on page 89, and how to register XML schema using a Java program in Example 7-4 on page 138. Refer to *DB2 10 for z/OS pureXML Guide*, SC19-2981 for more information.

Naming standard and convention for schemas

As we saw in Chapter 5, “Validating XML data” on page 73, the XML schema repository offers a wide range of choices for specifying XML schemas when validating XML documents against a schema, both for automatic and explicit validation.

You have the option of specifying a schema using

- ▶ Schema name
- ▶ URI and location hint
- ▶ Namespace

Although it is certainly convenient to be able to choose the option you prefer, it is probably a good idea to decide and document what method you want to use in your company. This should help ensure that it is always crystal clear which XML schema(s) you are working with, without having to go through the quite elaborate rule set for schema selection.

In addition, it might be worth while to define a naming standard for the XML schemas taking into account that XML schemas may evolve over time and require several versions available in the XSR at a given time.

Neither of these are technical tasks, but they often lie within the responsibilities of the DBA and may be good to plan for when starting a new XML project.

10.2.6 Creation of XML indexes

You can create an index on an XML column for efficient evaluation of Xpath expressions to improve performance of queries on XML documents. In contrast to simple relational indexes where index keys are composed of one or more table columns that you specify, an XML index uses a particular Xpath expression to index paths and values in XML documents stored in a single XML column.

You should specify a data type for every XML index. XML indexes support the data types VARCHAR, DECFLOAT, DATE, and TIMESTAMP. In Example 10-3 it is shown how to create an XML index.

Example 10-3 Create an XML index

```
CREATE INDEX IXMLNTRY
ON BK_TO_CSTMTR_STMT(BK_TO_CSTMTR_STMT)
GENERATE KEY USING XMLPATTERN
'declare default element namespace
  "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
  /Document/BkToCstmtrStmt/Stmt/Ntry/BookgDt/DtTm'
AS SQL TIMESTAMP
```

For performance considerations for XML indexes, reference 11.3, “Managing access path selection with XML” on page 249.

10.2.7 Grants and authorizations required

When a table is created with an XML column, an XML table space, XML table, and a NODEID index and document ID index are implicitly created. The privilege set must include the following privileges:

- ▶ The USE privilege on the buffer pool and the storage group that is used by the XML objects.
- ▶ If the base table space is explicitly created, CREATETS is also required on the database that contains the table (DSNDB04 if the database is implicitly created)

If you add an XML column, the privilege set requires the CREATETAB and CREATETS privileges on the database that contains the table (DSNDB04 if the database is implicitly created), and USE privilege on the buffer pool and the storage group that is used by the XML objects.

The implicitly created objects are owned by the owner of the base table.

10.3 Housekeeping

For XML objects maintenance, You can use IBM DB2 for z/OS utilities. The utilities handle XML objects similar to the way that they handle LOB objects.

- ▶ CHECK DATA

In addition to normal checking, the CHECK DATA utility also checks XML relationships, the integrity of XML documents and system-generated indexes that are associated with XML data.

- ▶ LOAD/UNLOAD

The input/output data can be in the textual XML format or the binary XML format.

If you load data into an XML column that has an XML type modifier, the LOAD utility validates the input data according to the XML schema that is specified in the XML type modifier.

Note that by using spanned recors, you are able to unload and load XML documents with a file size in excess of 32 KB with good performance. Note that the crossloader capability of the LOAD utility does not support the XML data type.

- ▶ REORG TABLESPACE

You can use the REORG TABLESPACE utility to reorganize XML objects.

When you run REORG on an XML table space that supports XML versions, REORG discards rows for versions of an XML document that are no longer needed.

- ▶ RUNSTATS

You can use the RUNSTATS utility to gather statistics for XML objects.

Note that in the physical implementation, an XML document may take up more than one row in the XML table space, so for the real time statistics of XML table spaces in SYSIBM.SYSTABLESPACESTATS, the number of rows is reported, not the number of XML documents.

For more information about utility support, please reference Chapter 9, “Utilities with XML” on page 181

10.4 Backup and recovery

Like a LOB column, an XML column holds only a descriptor of the column. The data is stored separately. Backup and recovery of XML objects is quite similar to backup and recovery of LOBs. A base table space must be kept consistent with its' associated LOB or XML table

spaces with respect to point-in-time recovery, so for backup and recovery, you should group all related objects together.

Use the REPORT utility with TABLESPACESET option to identify related objects which may include objects related by RI or auxiliary relationships to one or more XML and LOB table spaces. An example can be found at 9.14, “REPORT” on page 215.

You can also use LISTDEF with XML/LOB and RI option to include related objects as a list. Reference 9.6, “LISTDEF” on page 193 for more detail.

10.5 Diagnostics

When dealing with errors in XML table spaces and related objects, most of the problems are the same as with any other table spaces and the usual techniques can be applied to diagnose and solve the problems.

The following are errors that are specific to XML:

- ▶ Corrupted XML document(s). There may be missing rows in a document (which can be made up of more than one row) or structural defects to the nodes.
- ▶ Inconsistencies between XML table space and NODEID index. An index entry may exist but no corresponding XML document and vice versa.
- ▶ Inconsistencies between base table space and NODEID index. A reference may exist in the base table but no corresponding entry in the NODEID index, and vice versa.
- ▶ XML documents have not been validated against any schema in an XML type modifier.
- ▶ One or more XML documents are not valid according to any of the schemas in the XML type modifier.

10.5.1 Identification of XML related objects

There are several ways of determining the objects related to the XML column:

- ▶ Run the REPORT TABLESPACE utility to identify all objects belonging to base table space set and XML table space set. An example is shown in Example 9-20 on page 215.
- ▶ Query the DB2 catalog to obtain information of all the related objects. Various example queries are shown in 4.4, “Catalog queries to gather information” on page 63.
- ▶ Use LISTDEF with keyword ALL or XML in conjunction with the utilities to include all or all XML related index and/or table spaces. For more information, refer to 9.6, “LISTDEF” on page 193.

10.5.2 Investigating XML specific errors

In problem analysis one of the first actions is often to do a display of the table and index spaces in question to determine whether the XML or base table is in a restricted state.

The output of a display command is shown in Example 10-4. We see that all spaces are in status RW except the XML table space which is in the restricted state CHKP, check pending. The XML table space is easily identifiable with the type of XS.

Example 10-4 Display database command shows XML table space in AUXW

```
DSNT360I  -DB0B *****
```

```

DSNT361I  -DBOB *  DISPLAY DATABASE SUMMARY
              *  GLOBAL
DSNT360I  -DBOB *****
DSNT362I  -DBOB      DATABASE = DSN00155  STATUS = RW
              DBD LENGTH = 4028
DSNT397I  -DBOB
NAME      TYPE PART  STATUS              PHYERRLO PHYERRHI CATALOG  PIECE
-----
BKRTORCS TS      0001 RW
BKRTORCS TS              RW
XBKR0000 XS      0001 CHKP
XBKR0000 XS              CHKP
IRDOCIDB IX      L0001 RW
IRDOCIDB IX        L*  RW
IRNODEID IX      L0001 RW
IRNODEID IX        L*  RW
IXMLNTRY IX      L0001 RW
IXMLNTRY IX        L*  RW
***** DISPLAY OF DATABASE DSN00155 ENDED *****
DSN9022I  -DBOB DSNTDDIS 'DISPLAY DATABASE' NORMAL COMPLETION

```

There are only three restricted states that are related specifically to XML, and they may appear in other contexts too with another meaning. These are listed in Table 10-4.

Table 10-4 Restricted states related to XML

Restricted state	Cause	Comment
AUXW on base table	Invalidated XML column as a result of running CHECK DATA AUXERROR INVALIDATE	Can also be related to a LOB column of the table.
ACHKP on base table	Invalid XML column detected when running CHECK DATA AUXERROR REPORT	Can also be related to a LOB column of the table.
CHKP on XML table	Validation of documents have not been made though required because of change in XML type modifier	

Run the CHECK DATA utility with the INCLUDE XML TABLESPACES keyword to determine the exact cause of any of these states, and what further action to take. For a comprehensive description of the XML-related capabilities of the CHECK DATA utility, refer to 9.1, “CHECK DATA” on page 182.

Run the CHECK INDEX utility if you suspect errors in any of the indexes, like missing or extraneous entries. This is applicable to the DOCID index on the base table, and the NODEID index on the XML table, as well as any XML value indexes you have created yourself. For more information on the CHECK INDEX utility, refer to 9.2, “CHECK INDEX” on page 186.

Run the REPAIR utility with the LOCATE KEY and LOCATE RID keywords to ascertain whether certain rows and/or index entries exist. LOCATE KEY can be used to locate a row in a base table using the DOCID key in the DOCID index. LOCATE RID can be used to locate a row in the XML table space using the RID of the row. Example 10-5 shows how to use the

utility to locate a row in base table space BKRTORCS with DocID =2, and a row in the XML table space with RID X'00000201'.

Example 10-5 REPAIR LOCATE control statements for diagnosing XML inconsistencies

```
REPAIR LOCATE TABLESPACE DSN00155.BKRTORCS
  KEY(2) INDEX XMLR2.I_DOCIDBK_TO_CSTMR_STMT DUMP
REPAIR LOCATE TABLESPACE DSN00155.XBKRO000
  RID X'00000201' DUMP
```

Note: Currently open APAR PM26592 should be applied to allow for use of BIGINT data type in REPAIR LOCATE KEY.

10.5.3 Correcting XML data

Depending on the type of error or inconsistency identified, there are different ways of correcting the data.

If the errors were identified by the CHECK DATA utility, you can use the CHECK DATA utility to correct the problems.

- ▶ Run CHECK DATA with SHRLEVEL REFERENCE and XMLERROR INVALIDATE. This will cause the utility to
 - delete invalid XML documents and move them to exception tables
 - invalidate XML entries in the base table
- ▶ Run CHECK DATA with SHRLEVEL CHANGE and XMLERROR INVALIDATE. This will cause the utility to generate control statements for you to execute, including
 - REPAIR LOCATE DOCID DELETE statements for deleting orphan rows
 - REPAIR LOCATE RID REPLACE statements for invalidating entries in the base table
 - REBUILD INDEX for the NODEID index if this is in error.

For details on the CHECK DATA utility and the options for correcting the data, refer to 9.1, “CHECK DATA” on page 182.

Notes:

PTF UK62510 for APAR PM24947 should be applied for the CHECK DATA SHRLEVEL CHANGE utility to generate REPAIR statements.

APAR PM21834 (currently open) provides various usability fixes for DB2 utilities including CHECK DATA.

If the errors were identified by the CHECK INDEX utility, possibly in conjunction with REPAIR LOCATE, the appropriate course of action is shown in Table 10-5.

Table 10-5 Corrective action after running CHECK INDEX

Problem	Solution
Error in DocID index	Ensure table space is at correct level Rebuild index

Problem	Solution
Mismatch between NODEID index or user-defined XML index and XML table space, and index is correct.	Use REPAIR LOCATE RID DELETE to remove orphan row
Mismatch between NODEID index or user-defined XML index and XML table space, and XML table space is correct	Rebuild index

Finally, you may need to reset the restrictive status by running CHECK DATA SHRLEVEL REFERENCE, or by using the REPAIR utility as shown in Example 10-6.

Example 10-6 Using REPAIR utility to clear ACHKP status on table space

```
REPAIR OBJECT
SET TABLESPACE DSN00155.BKRTORCS NOAUXCHKP
```



Performance considerations

In this chapter we provide a checklist of the major (additional) performance considerations when deploying an application that uses pureXML. The checklist covers

- ▶ Choice of relational or XML storage
- ▶ XML Schema validation
- ▶ Managing access path selection with XML
- ▶ Encourage use of native SQL DB2 routines
- ▶ External language programming
- ▶ DBA considerations
- ▶ SQL/XML coding techniques

11.1 Choice of relational or XML storage

The very first thing to consider is whether native XML is the right choice of storage model for persisting the data in your new application. DB2 offers three choices:

- ▶ Relational-only storage (which we will not discuss further in this chapter)
- ▶ XML only storage
- ▶ Hybrid storage

11.1.1 XML only storage

The option to store data using tables that only contain columns with the XML data type, is a totally practical option. The enforcement of XML schemas, combined with the ability to create XML indexes to enforce uniqueness and XML indexes for performance provides the ability to implement an XML database with integrity and performance. Figure 11-1 illustrates two ISO20022 message types, covering the area of payment notifications, implemented as an XML-only database in DB2.

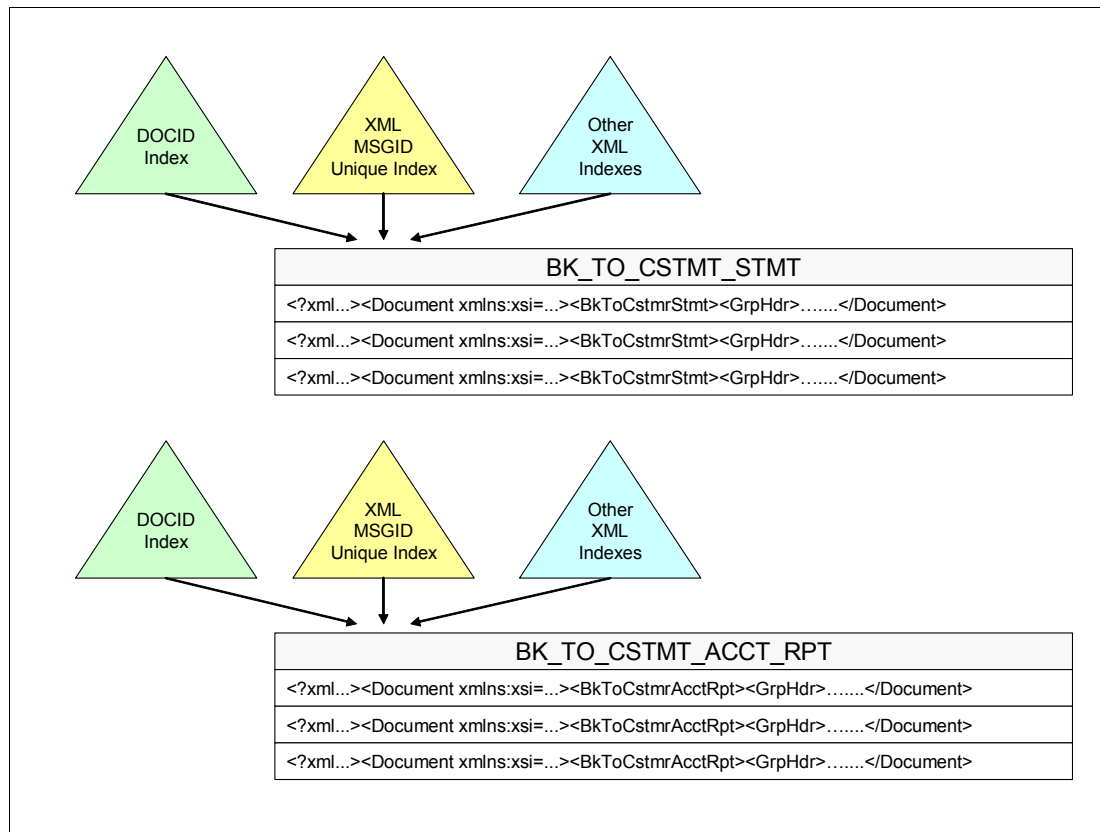


Figure 11-1 XML-only database design

An XML-only design has a number of potential drawbacks for traditional DB2 users.

1. You cannot define referential integrity constraints between tables based on XML columns, or XML expressions against the contents of those XML columns.
2. Development and DBA teams would necessarily be forced to **think** in terms of XML and XPath expressions for every database interaction that they performed. This would be a large conceptual shift from traditional relational and XML thinking.

XML-only storage would be a radical change for many DB2 users, which might stretch existing skills beyond their comfort zones. For most DB2 users, some kind of hybrid storage model will be best, because it will minimise the learning curve for DBAs and developers who are established relational users, without diminishing the ability to manage XML documents with pureXML.

11.1.2 Hybrid storage

The ability of DB2 to support a hybrid storage model, as illustrated in Figure 11-2, provides the best of both worlds to DB2 users. The XML documents can be preserved in their original state (which may be needed for compliance reasons in some cases) and can be accessed at any level of detail using SQL/XML. Optionally, those data elements that benefit from being stored in relational format can be stripped out and stored in relational columns. (But don't strip XML data elements out into relational storage, unless there is a clear benefit in doing so).

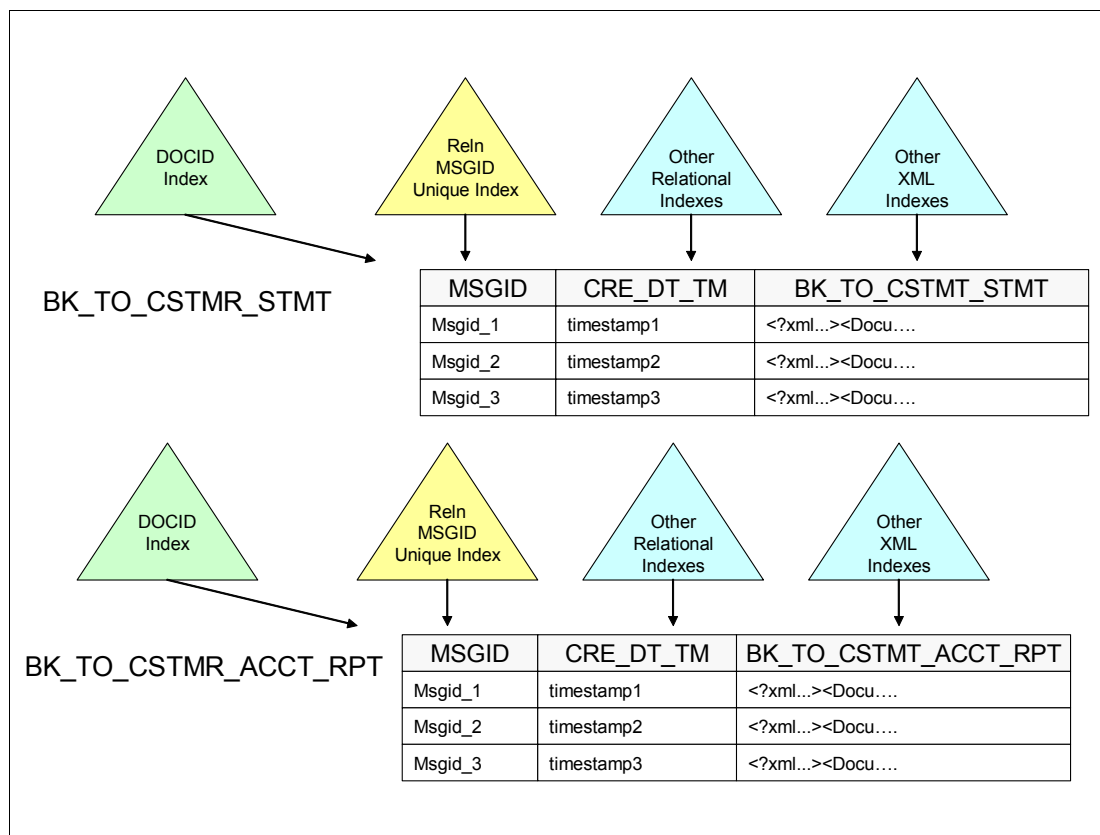


Figure 11-2 Hybrid storage model

The programming examples in this book (SQL procedures, COBOL and Java) have all shown how easy it is to use the SQL/XML language to strip out data elements of interest, and store them as relational columns. While this is not necessary to achieve a high performance database, it can be desirable for many other reasons.

- ▶ Referential Integrity can be defined on the indexed relational columns.
- ▶ The database is easier to work with, because it can be represented easily in the familiar representation of Entity-Relationship diagrams, which is the conceptual bedrock upon which traditional relational systems have been designed for years.
- ▶ The contents of the XML documents can be accessed with ease, using the power of the SQL/XML language.

- ▶ Developers need not become super-skilled in XML expressions, particularly if the database is implemented with appropriate use of stored procedures and user defined functions (to encapsulate XML functions in callable SQL routines). This will allow the developers to work very effectively with the database using “plain old SQL” for much of the time.
- ▶ Data structures that are naturally best supported in an XML structure can be stored unaltered in DB2 in their optimal physical representation, and can be accessed efficiently using XML expressions in SQL/XML.
- ▶ The inherent strengths of the XML structures can still be fully utilized with SQL/XML, regardless of the fact that we have also stored some of the data elements in relational columns to make some development and DBA tasks easier.
- ▶ Relational and XML data can be joined and processed together with ease and good performance. New applications can be written to fully utilise pureXML, while being able to integrate seamlessly with traditional relational structures.
- ▶ Data that is very stable in terms of its structure (low variability) and accessed very often can be a good fit for extraction into relational columns. More variable parts of the data are better kept in XML.

Dedicated use of XML columns

Another database design consideration is whether it's better to have one XML column to contain XML documents of differing XML schemas, or multiple columns that each contain documents of one specific XML schema.

Either approach is valid. This is a trade-off between database management complexity and performance.

If you have multiple XML columns to store XML documents of different XML schemas, then you will create many more database objects (table spaces, data sets, DOCID indexes, NODEID indexes, etc....). The logical data model is easier to understand if different XML document types are stored in their columns. However, the physical data model will have a larger number of objects to track.

XML Indexes and performance may suffer if you combine multiple XML document types into a single DB2 column. If you create an XML index on an XML column with multiple XML schemas then the costs of index creation will increase, even if many of the documents do not contain matching nodes. Additionally, if you are not careful about XPath expressions and namespaces, you might get unexpected results if you get an index hit against an XML document type that you were not targeting in your application.

As a rule of thumb, it is usually best to store XML documents in DB2 columns that are dedicated to documents of that particular XML schema.

11.1.3 Natural fit for XML storage

Some data are an obvious natural fit for pureXML storage. The ISO20022 standard for banking messages used in the application scenarios of this book is a good example. These messages are defined by ISO as a standard for messaging for european banking. It will be important to store these messages in exactly the way that they were transmitted, because they are a record of financial transactions. DB2 pureXML allows us to store these messages in their original format, and also allows us to query and interrogate them using the power of SQL/XML, and the efficiency of pureXML.

There are a large range of other industry standard models for XML storage and XML messaging, particularly in financial services, but also in other industries. We may wish to store in a persistent database for both query and transactional purposes. DB2 pureXML provides a good foundation for such systems.

Many data structures for in-house systems would be best implemented using XML. The example of insurance quotations comes to mind. The range of data collected for a car insurance proposal can vary dramatically depending on the vehicle, the applicants, the usage of the vehicle, the claims history and so on. Do we really want to design a relational database with 50 or so tables to store millions of internet quotes, when they are sparsely populated and only 1% of them will be taken forward into a purchased policy? Surely this is an example of an in-house data structure that should implemented in XML. Price comparison sites must surely use XML documents and web services to communicate between different insurers. With pureXML these documents could be stored for efficient retrieval, and indexed and queried in their native format for patterns in client behavior.

The following is a list of the data characteristics that tend to make XML a good choice for data storage.

- ▶ Hierarchical data

The XML data model is naturally hierarchical, which means that it will tend to be a good fit for data that is naturally hierarchical.

- ▶ Semi-structured data

XML schemas can be rigid or loose as appropriate. They can have rigid constraints for some XML nodes, and looser constraints for other XML nodes as appropriate, to accommodate the appropriate level of structure that is required.

- ▶ Document/narrative data

XML is always accused of being verbose!

- ▶ Many different schemas

The ISO20022 standard is a case in point. It contains hundreds of schemas. The development cost of building relational database schemas for all these message structures would be astronomical. With pureXML the number of tables and columns could be very small, because the data structures are managed in the XML schemas, with very little DBA impact.

- ▶ Large schemas (with sparsely populated attributes/elements)

The ISO20022 schemas are very large to provide a generic schema that could cater for a very wide range of payments scenarios. The average personal banking customer will only use a subset of the facilities in that schema. No additional processing or storage costs are incurred for leaving the majority of the XML nodes empty of data. XML indexes will only contain entries for those XML patterns where an indexed element or attribute is found, thus minimizing the space required for XML indexes.

- ▶ Quickly evolving schema

If a relational database schema changes, it can take months for the development, testing and QA processes to introduce a new release of the application. A major part of that release time is taken up by database change management. This part would be substantially reduced if all you had to do was add a new schema document to the existing XSR schema, to define a new release of the XML schema. Furthermore, database schema changes do not necessarily require data migration: the application can decide whether an XML document which is compliant with the old schema release needs to be modified to accommodate the new schema release.

- ▶ Forms-based applications

If the data from a forms based application is created in XML, why would you want to convert it to relational for storage?

- ▶ Data with nulls and multiple values

XML storage is a way for DB2 to support multi-value data storage.

- ▶ Existing industry standard schemas for XML

The ISO20022 examples in this book have shown how easy it is to incorporate an existing XML standard into DB2, and use it with a minimal amount of database administration work. There are many other industry-related XML standards, some of which are listed in chapter 1.

If the data for your application fits some or all of these characteristics, then pureXML could be the best physical storage option. If you choose pureXML storage, then consider carefully the optimal hybrid storage design, which makes development tasks easy, and allows the strengths of the XML model to be fully utilized.

11.2 XML Schema validation

Having chosen pureXML for your storage model, and the appropriate degree of hybrid storage, the next questions are whether and when to perform schema validation. The primary considerations here are

- ▶ Data integrity is paramount. XML schemas are the method by which the integrity rules of XML data are defined, and XML schema validation is how they are enforced.
- ▶ Schema validation is CPU intensive, so we don't want to validate XML documents unnecessarily or repeatedly.
- ▶ XML schema validation is a candidate for zIIP and zAAP redirect, which means that XML schema validation need not increase your general purpose CPU resource consumption.

Some initial guidance follows:

- ▶ Don't take chances with data integrity. If you are not sure that an XML document is valid (perhaps if it was received from an external source, which will be a common scenario), then you should validate it before you commit it to your database.
- ▶ If you know that an XML document has been validated against its schema (perhaps because it was validated in WebSphere Message Broker before being stored in DB2) then do not re-validate it just for the sake of it.
- ▶ In scenarios where you will always want to perform schema validation, ensure that the DB2 table is defined with an XML type modifier, so that XML validation cannot be bypassed.
- ▶ In scenarios where the some, but not all, of the XML documents have been validated before they are presented to DB2, do not define the DB2 table with an XML type modifier. In this case, make use of the DSN_XMLVALIDATE function to perform validation only where necessary.
- ▶ You can optionally control whether certain users or programs need to perform XML validation by restricting access to the base DB2 table, and encapsulating write access (with or without XML schema validation calls) to the DB2 table within stored procedures or user defined functions. Access to these routines and functions can be granted to controlled groups as appropriate.
- ▶ You could implement automated XML validation in test and quality assurance environments in order to flush out XML data quality issues before an application is

deployed to the production environment. The amount of XML validation that is enforced in a production environment with higher transaction volumes could be set to a lower level if the XML documents in a production environment are known to be valid.

It is worth noting that the XML validation process (whether invoked manually via `DSN_XMLVALIDATE`, or automatically via DDL table modifier) always requires a string data type as input. If you invoke XML schema validation against an XML data type, then the validation first converts the XML document to a string value, and then performs XML schema validation. Validation using binary XML is not yet supported.

You should also monitor zIIP and zAAP utilization, because XML processing is 100% eligible. A discussion paper about zIIP and zAAP monitoring is available at:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101227>

11.3 Managing access path selection with XML

XML indexes can be used instead of relational indexes, or in addition to relational indexes. The critical consideration is that DB2 indexes (of one sort or the other) are used to eliminate the retrieval and searching of a large number of XML documents. If scanning a relational table space is bad, then the prospect of scanning and processing an XML table space as well would be far worse.

The principles of designing access paths are no different from relational access path selection.

- ▶ XML indexes are B-tree structures just like relational indexes
- ▶ The DB2 Optimizer will consider using both XML and relational indexes based on their attractiveness
- ▶ The attractiveness of XML indexes is based on largely the same measures as relational indexes. (cardinality, filter factors etc... that you will find in the DB2 catalog)

The best way of approaching the subject of XML indexes for access path selection is to be aware of the differences that exist between XML index usage and relational index usage, and add this knowledge to traditional quality assurance processes that you already use for access path selection.

11.3.1 Differences between XML and relational indexes

The purpose of an XML index is to retrieve a number of DOCID values, which will be used to perform filtering of qualifying rows in a table as shown in Figure 11-3. The subsequent table access will be performed via the index on hidden `DB2_GENERATED_DOCID_FOR_XML` column.

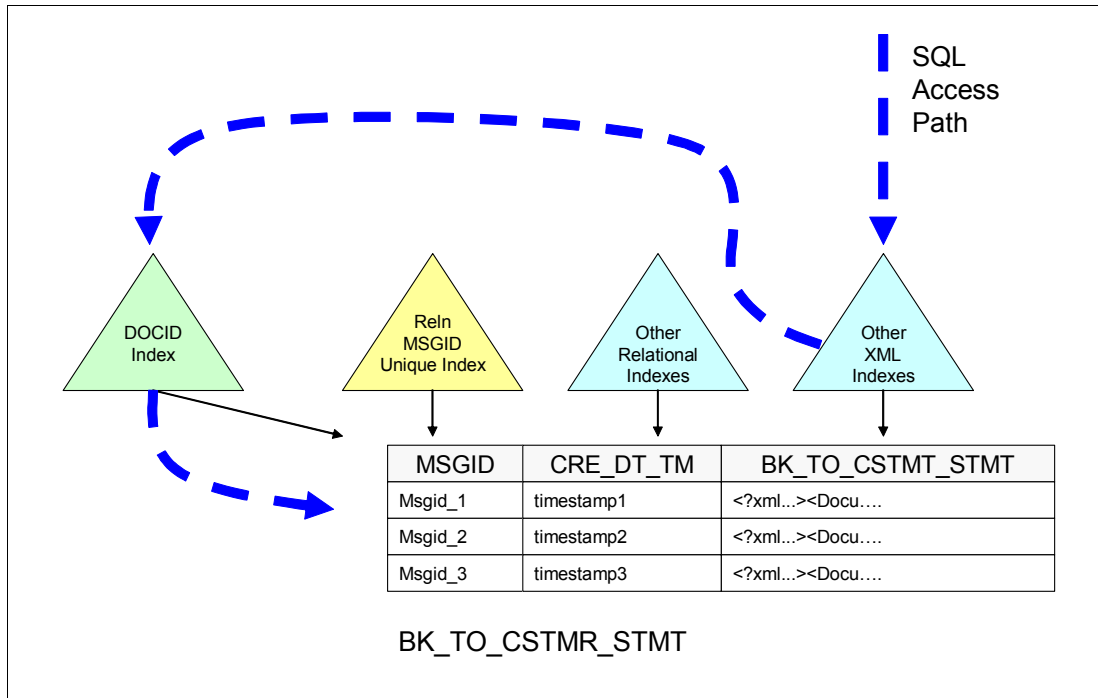


Figure 11-3 Physical access path using an XML Index.

XML indexes are created on a XML Pattern (examples of XML patterns are provided shortly). The result of that XML pattern yields a result (which may be very variable in nature, depending on the content of the XML document and the constraints of an XML schema).

The differences between XML and relational indexes are:

- ▶ Relational indexes may be defined on one or more relational columns. XML Indexes can only be defined on one XML element or attribute (using an XML pattern expression)
- ▶ Relational indexes always have one index entry for every row in a table. However XML indexes are much less prescriptive. XML indexes are based on an XML pattern. An XML pattern may occur any number of times in an XML document. So, XML Indexes may contain 0, 1 or many entries for each row in the table.
- ▶ Relational indexes are always based on the data types of the column(s) that they are defined on. The data types found at the locations of an XML pattern may be many and varied unless an appropriate XML schema is enforced. XML indexes are defined based on a mapping to a particular data type, but whether or not the data type that is found at that location can be cast to that data type for an index match to be achieved depends on the degree to which the XML schema constrains the data contents.
- ▶ Relational Indexes can be used to support table clustering. XML indexes may not be used for table clustering support.

Having understood the nature of XML indexes and how they differ from relational indexes, we must now use that understanding to consider how to design our XML indexes.

11.3.2 XML index design

The differences between XML indexes and relational indexes lead to a set of XML index design considerations which follow:

XML Index patterns

A wide range of XML patterns can be used in XML indexes.

At the most restrictive end of the spectrum are “lean” indexes. The index in Example 11-1 is targeted at a data element in a specific XPath location, which is cast to a specific relational data type. In this example an index is built based on mapping the data elements at XPath /Document/BkToCstmrStmt/GrpHdr/MsgId and casting whatever is found there to Varchar(35). Note: XMLPATTERNS must define the appropriate namespace if the Documents contain a namespace declarations.

Example 11-1 A lean XML index

```

Create index LEAN_XMLIX
  ON BK_TO_CSTMRT_STMT(BK_TO_CSTMRT_STMT)
  Generate Key using XMLPATTERN 'declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    /Document/BkToCstmrStmt/GrpHdr/MsgId'
as SQL VARCHAR(35) ;

```

At the least restrictive end of the spectrum are “heavy” indexes. The index in Example 11-2 is targeted at any occurrence of a data element called MsgId within the entire document.

Example 11-2 A heavy XML index

```

Create index HEAVY_XMLIX
  ON BK_TO_CSTMRT_STMT(BK_TO_CSTMRT_STMT)
  Generate Key using XMLPATTERN 'declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    //MsgId'
as SQL VARCHAR(35) ;

```

Choosing lean or heavy XML indexes is a physical design trade off. You want to maximise the filter factor of every index, but you might be prepared to compromise the filter factor if you can get away with a smaller number of indexes.

When you create your first XML index, you will want to get some confirmation about whether it is doing the task that you want it to do. The obvious way to do this is to explain a query that should benefit from the index, and see if the index is selected by the optimizer.

Consider the “silly” XML index in Example 11-3. This index was created with an XMLPATTERN that did not actually match any nodes in the XML document, because of a typographical error in the xpath expression.

Example 11-3 A “silly” XML index

```

Create index SILLY_XMLIX
  ON BK_TO_CSTMRT_STMT(BK_TO_CSTMRT_STMT)
  Generate Key using XMLPATTERN 'declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    /Document/BkToCstmrStmt/MsgId'
as SQL VARCHAR(35) ;

```

If you find that the index is not selected, and you can't figure out why the optimizer is not choosing the index, then you should take a moment to check the contents of SYSIBM.SYSINDEXES. The table snapshot in Example 11-4 shows a subset of the catalog

statistics for the lean and the heavy indexes that were created in example and example. It also shows the catalog statistics of another XML index which is based on an XML pattern with a typographic error.

- ▶ The table has 5 rows, each with an XML document
- ▶ The lean XML index has a cardinality of 5, because it is based on an XPath expression that identifies the MsgId data element, which happens to be unique
- ▶ The heavy XML index has a cardinality of 24, because the MsgId data element is repeated multiple times in a typical Bk_To_Cstmr_Stmt message
- ▶ The “silly” index has a cardinality of 0!

```

-----+-----+-----+-----+-----+-----+-----+-----+-----
select
  substr(creator,1,5) concat '.' concat substr(name,1,23) as IndexName,
  int(firstkeycardf) as firstkeycardf,
  int(fullkeycardf) as fullkeycardref
  from sysibm.sysindexes
  where creator = 'XMLR3' and
  tname = 'BK_TO_CSTMR_STMT'

-- yields

-----+-----+-----+-----+-----+-----+-----+-----+-----
INDEXNAME                                FIRSTKEYCARDF  FULLKEYCARDREF
-----+-----+-----+-----+-----+-----+-----+-----+-----
XMLR3.HEAVY_XMLIX                        6              24
XMLR3.I_DOCIDBK_TO_CSTMR_STMT            5              5
XMLR3.LEAN_XMLIX                         5              5
XMLR3.RELN_IX1                           5              5
XMLR3.SILLY_XMLIX                        0              0
DSNE610I NUMBER OF ROWS DISPLAYED IS 5
DSNE616I STATEMENT EXECUTION WAS SUCCESSFUL, SQLCODE IS 100
-----+-----+-----+-----+-----+-----+-----+-----+-----

```

Figure 11-4 Catalog query to sysibm.sysindexes

The fact that FIRSTKEYCARDF and FULLKEYCARDF are zero is the warning alarm that tells you that the index did not point to any matches at all in any of the XML documents in the table. When this happens, you can deduce that you have an error in your XML pattern in the index definition. Check the XML pattern, correct it, recreate the index, run RUNSTATS and review the cardinality statistics until you get a sensible number.

This is a very useful way to check that the index you have created is based on an XML pattern that finds hits in the data. This example illustrates one of the differences between relational indexes and XML indexes very well. It is impossible to create a relation index that doesn't have a pointer to every row in the table. But it is very possible to have a perfectly valid XML index that points to zero rows in the table.

As an aside, there is a learning curve for the SQL programmer to become comfortable with writing XML expressions such as XMLTABLE, XMLQUERY and XMLEXISTS. If you code an invalid SQL/XML statement you will receive an SQL error (such as SQLCODE -104) which will usually contain some helpful guidance as to what is wrong with your statement. What can be more frustrating is when you code an SQL/XML statement that is valid, but it returns no rows. The two most common causes of getting no rows are XPath expressions containing a

typographic error, and incorrect namespace declarations. These two reasons are also the first things you should look at when an XML index matches nothing in the table.

XML index maintenance cost

XML indexes are more expensive for insert than relational indexes because DB2 has more work to perform. Consider inserting a row into a DB2 table with an integer index and an XML index. Ignoring factors like page splits, the maintenance of the relational index is just a case of inserting a new entry into the appropriate index leaf page. However, the maintenance of the XML index requires that the XML document is scanned for zero, one or many matches to the `xmlpattern` of the index, as well as the insert of a new value into the index leaf page. The maintenance of an XML index is similar to the maintenance of an index-on-expression (with the XPath XMLPATTERN being the *expression*).

XML indexes are also more expensive for select than relational indexes because two indexes must be used:

- ▶ The NODEID index on the XML table space must be used to find matches for the XML pattern that is being searched on
- ▶ The DOCID index must be used next, to retrieve RIDs to perform table access.

XML indexes should therefore be chosen with care, because they will increase the path length of insert operations more than relational indexes. This statement should NOT be taken to mean that all searchable data fields should be stripped from XML documents and placed in indexed DB2 columns. The fact that an XML index adds some extra path length must be balanced against the flexibility of being able to index directly into any part of a large XML document without having to maintain a copy of that data element in a separate relational column.

XML index eligibility

There are a number of index eligibility constraints that you need to be aware of when designing XML indexes. When considering index eligibility, it is always useful to remember that the purpose of an XML index is to return a small (hopefully) number of DOCID values which will be used for table access.

XMLEXISTS functions are eligible for index access, provided the XMLEXISTS predicate is supported by the XML pattern of the XML index. Example 11-4 shows an SQL/XML query with an XMLEXISTS predicate that could be supported by either the lean or the heavy XML indexes. In this case, the lean XML index was chosen, as proven by the results of an explain request.

Example 11-4 Explain for XMLEXISTS

```
explain plan set queryno = 99 for
  select c.msg_id, c.msg_cre_dt_tm
         from BK_TO_CSTMR_STMT c
         where   xmlexists('declare default element namespace
                          "urn:iso:std:iso:2002:tech:xsd:camt.053.001.02";
                          $i/Document/BkToCstmrStmt/GrpHdr[MsgId="AAAASESS-FP-STAT002"]'
                          passing c.BK_TO_CSTMR_STMT as "i");

select
  planno, creator, tname, accesstype,
  matchcols, accesscreator, accessname, table_type
  from XMLR3.PLAN_TABLE where queryno = 99;
```

```
PLANNO          = 1
CREATOR         = XMLR3
```

```

TNAME          = BK_TO_CSTMR_STMT
ACCESSTYPE     = DX
MATCHCOLS      = 1
ACCESSCREATOR  = XMLR3
ACCESSNAME     = LEAN_XMLIX
TABLE_TYPE     = T

```

XMLTABLE without a filtering predicate is not generally eligible for XML index access. This is because the result of an XMLTABLE operation is just a table without any filtering of DOCIDs, unless a predicate is included alongside the XMLTABLE function. Example 11-5 shows the use of the XMLTABLE function in conjunction with an XMLEXISTS predicate. The access path is to use the lean XML index to retrieve the DOCIDs, and then to use the table function X (which is the name assigned to the result of the XMLTABLE function) to retrieve the data values.

Example 11-5 Explain for XMLTABLE with XMLEXISTS

```

EXPLAIN PLAN SET QUERYNO = 88 FOR
  SELECT X.MSG_ID, X.CRE_DT_TM FROM BK_TO_CSTMR_STMT as C,
    XMLTable(XMLNAMESPACES(DEFAULT
      'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02',
      '$d/Document/BkToCstmrStmt'
      PASSING c.BK_TO_CSTMR_STMT as "d"
      COLUMNS
        "MSG_ID" VARCHAR(35) PATH './GrpHdr/MsgId/text()',
        "CRE_DT_TM" TIMESTAMP PATH './GrpHdr/CreDtTm/text()' ) AS X
  where xmlexists('declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $i/Document/BkToCstmrStmt/GrpHdr[MsgId="AAAASESS-FP-STAT002"]'
    passing c.BK_TO_CSTMR_STMT as "i");

```

```

select
  planno, creator, tname, accesstype, matchcols,
  accesscreator, accessname, table_type
  from XMLR3.PLAN_TABLE where queryno = 88;

```

```

PLANNO        = 1
CREATOR       = XMLR3
TNAME         = BK_TO_CSTMR_STMT
ACCESSTYPE    = DX
MATCHCOLS     = 1
ACCESSCREATOR = XMLR3
ACCESSNAME    = LEAN_XMLIX
TABLE_TYPE    = T

```

```

PLANNO        = 2
CREATOR       = XMLR3
TNAME         = X
ACCESSTYPE    = R
MATCHCOLS     = 0
ACCESSCREATOR =
ACCESSNAME    =
TABLE_TYPE    = F

```

Another way of getting XMLTABLE to use an XML Index does not require the use of XMLEXISTS. You can specify a predicate on an XMLTABLE function, as shown in Example 11-6, and get XML index access.

Example 11-6 Explain for XMLTABLE with an XML predicate

```
EXPLAIN PLAN SET QUERYNO = 55 FOR
  SELECT X.MSG_ID, X.CRE_DT_TM
     FROM BK_TO_CSTMR_STMT as C,
     XMLTable(XMLNAMESPACES(DEFAULT
       'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02',
       'http://www.w3.org/2001/XMLSchema-instance' AS "xsi"),
       '$d/Document/BkToCstmrStmt[GrpHdr/MsgId="AAAASESS-FP-STAT002"]'
       PASSING c.BK_TO_CSTMR_STMT as "d"
     COLUMNS
       "MSG_ID"    VARCHAR(35)  PATH './GrpHdr/MsgId',
       "CRE_DT_TM"  TIMESTAMP    PATH './GrpHdr/CreDtTm/text()' ) AS X

select
  planno, creator, tname, accesstype, matchcols,
  accesscreator, accessname, table_type
  from XMLR3.PLAN_TABLE where queryno = 55;
```

```
PLANNO          = 1
CREATOR         = XMLR3
TNAME          = BK_TO_CSTMR_STMT
ACCESSTYPE     = DX
MATCHCOLS      = 1
ACCESSCREATOR  = XMLR3
ACCESSNAME     = LEAN_XMLIX
TABLE_TYPE     = T
```

```
PLANNO          = 2
CREATOR         = XMLR3
TNAME          = X
ACCESSTYPE     = R
MATCHCOLS      = 0
ACCESSCREATOR  =
ACCESSNAME     =
TABLE_TYPE     = F
```

XMLQUERY is not eligible for XML index access, because the purpose of XMLQUERY is to return an XML document. XMLQUERY is never used to filter the rows in a table. Hence, XMLQUERY should in general be used in conjunction with other relational or XML predicates that will be used for filtering, as illustrated in Example 11-7.

Example 11-7 Explain for XMLQUERY with XMLEXISTS

```
EXPLAIN PLAN SET QUERYNO = 77 FOR
  select xmlquery('declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $d/Document/BkToCstmrStmt/Stmt' passing BK_TO_CSTMR_STMT as "d")
     from BK_TO_CSTMR_STMT C
  where xmlexists('declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
```

```
$i/Document/BkToCstmrStmt/GrpHdr[MsgId="AAAASESS-FP-STAT002"] '
passing c.BK_TO_CSTMRTMT as "i");
```

```
select
  planno, creator, tname, accesstype, matchcols,
  accesscreator, accessname, table_type
from XMLR3.PLAN_TABLE where queryno = 77;
```

```
PLANNO          = 1
CREATOR         = XMLR3
TNAME           = BK_TO_CSTMRTMT
ACCESSTYPE     = DX
MATCHCOLS      = 1
ACCESSCREATOR  = XMLR3
ACCESSNAME     = LEAN_XMLIX
TABLE_TYPE     = T
```

11.4 Encourage use of native SQL DB2 routines

Native SQL stored procedures are probably the second most important DB2 device (after XML indexes) to ensure good performance and efficiency in your pureXML applications.

Native SQL procedures in DB2 are an incredibly good environment for developing high performance applications, with or without XML. They are well covered in *DB2 9 for z/OS Stored Procedures: Through the CALL and Beyond*, SG24-7604, which explains some of their strengths as follows:

- ▶ Multiple procedural steps can be executed within DB2, eliminating network delays that an external program would experience making each SQL call and waiting for the results.
- ▶ Stored procedures run entirely within the DB2 engine.
- ▶ Procedural statements are converted to a native representation that is stored in the DB2 catalog and directory, as it is done with other SQL statements.
- ▶ zIIP eligible if they are called through DRDA with TCP/IP, type 4 Java

The advent of pureXML makes native stored procedures even more attractive for the following reasons:

- ▶ You can pass huge arrays of data to and from native stored procedures, via an XML IN or OUT data type
- ▶ You can avoid code page translations for XML processing. A stored procedure can retrieve an XML document from a DB2 table and work with it, entirely within the DB2 engine. By contrast, programming environments like COBOL that typically work in EBCDIC will need to convert the contents of the UTF-8 XML document into EBCDIC, in order to work with the data.
- ▶ XML procedures can receive an external XML document and parse it once, and then process it with SQL/XML many times without reparsing it.

In addition to the performance benefits of native stored procedures for XML processing, there are considerable development productivity benefits to be gained from creating frequently used procedures that encapsulate XML processing, and can be called by a developer using “plain old SQL”.

Native stored procedures provide a high performance platform for the programming of SQL and SQL/XML routines, which can be called from any application environment. Any application development project using DB2 z/OS should be evaluating which programming functions should be implemented as native stored procedures, so that they can be used and shared by all applications and application environments that use DB2.

11.5 External language programming

The biggest performance consideration for using external programming languages is the need to perform codepage translations. Chapter 8, “Using XML with COBOL” on page 155 discusses the fact that COBOL programs will be based in an EBCDIC environment, whereas the XML is stored in UTF-8. Therefore if you wish to exchange data between COBOL and pureXML in DB2, you must perform code page translations. Chapter 8 discusses the options to minimise the amount of code page translation that is performed.

The code page challenge is reduced for Java when using the binary XML format with the IBM Data Server Driver for JDBC and SQLJ, which presents binary XML data to the application only through the XML object interfaces.

11.6 DBA considerations

There are a number of database administration considerations that can have a significant impact on the performance of pureXML applications.

Use DB2 10 NFM universal table spaces

XML versioning support allows XML document updates to be performed with a high level of concurrency, but it depends on the underlying table spaces being a universal table space created in DB2 10 new function mode.

XML versioning is the mechanism that allows concurrent read access to an XML document whilst another user is updating it. It works by maintaining multiple copies of the XML document in the XML auxiliary table. It depends on the columns in the XML auxiliary table that are created in DB2 10 NFM. You are disabled from querying the XML table space directly, but Example 11-8 shows a catalog query and result which shows the columns that get created automatically in DB2 10 NFM.

Example 11-8 sysibm.syscolumns contents for auxiliary XML table space

```
select name, colno, coltype, length
  from sysibm.syscolumns
     where tbcreator = user
     and tbname = 'XBK_TO_CSTMR_STMT'
     order by colno ;
```

... yields

NAME	COLNO	COLTYPE	LENGTH
-----	-----	-----	-----
DOCID	1	BIGINT	8
MIN_NODEID	2	VARBIN	128
XMLDATA	3	VARBIN	15850
START_TS	4	BINARY	8

The XML table is structured as follows.

- ▶ The XML document (minus tags, which are stored in sysibm.sysstrings) is stored in the XMLDATA column, in one or more rows, depending on the size of the document.
- ▶ If the XML document is split into multiple rows, the minimum XML node id is stored against each row, and is indexed by the node_id index
- ▶ The DOCID is also stored in every row
- ▶ The start and end timestamps for each row

The DB2 10 "multi-versioning" table space format depends on the use of a cleanup SRB to remove old versions of XML docUMENTs created by UPDATE and DELETE. It's a small overhead, but this can result in some CPU activity and BP16K and index BP activity even when XML applications are idle.

Do not use DB2_GENERATED_DOCID_FOR_XML as a key

The DB2_GENERATED_DOCID_FOR_XML is not explicitly hidden, and is populated with a unique sequence number, and indexed, which makes it tempting to use as a primary key.

This column is an internal object, and IBM does not guarantee to keep it unchanged in the future. You should not develop applications that depend on this column.

Data compression

The storage of XML documents receives a certain amount of compression automatically. The XML tags are stored in SYSIBM.SYSSTRINGS and replaced with binary values. However, the contents of the data values in XML documents are not compressed by default.

The textual nature of many XML documents means that they are particularly well suited to compression techniques, which will lead to performance gains from reduced I/O and efficient use of buffer pool. You should therefore consider compressing DB2 tables with XML, and DB2 XML indexes.

After you specify COMPRESS YES in your DDL, the table space will be compressed at REORG time or with the new online compression available with DB2 10 during insert type operations.

The DSN1COMP utility can be used to assess the compression benefit, as usual.

Remember to REORG

You may not be allowed to view the contents of XML table spaces directly, but you are still responsible for reorganizing them.

The XML table space follows the same partitioning scheme as the base table space, but has the potential to grow much faster. Make sure that you REORG the XML table spaces frequently enough to maintain performance.

The method of reorganizing an XML table space requires that you build a REORG job for the base table space, and then you add an additional REORG control statement for each of the XML table spaces. The table spaces for the BK_TO_CSTMR_STMT example table are reorganized using the JCL in Example 11-9. Note that you must also specify the WORKDDN keyword on the REORG for the XML table space and provide the specified temporary work file. The default is SYSUT1.

Example 11-9 REORG of a table space with an XML table space

```

//REORG1 EXEC DSNUPROC,SYSTEM=DBOB,
//          LIB='DBOBT.SDSNLOAD',
//          UID=' '
//DSNUPROC.SYSPUNCH DD DSN=XMLR3.DBOB.CNTL.XMLR3DB.TSAUDIT1,
//          DISP=(MOD,CATLG),
//          SPACE=(TRK,(5,5),RLSE),
//          UNIT=SYSDA
//DSNUPROC.SYSREC DD DSN=XMLR3.DBOB.UNLD.XMLR3DB.TSAUDIT1,
//          DISP=(MOD,CATLG),
//          SPACE=(TRK,(15,5),RLSE),
//          UNIT=SYSDA
//DSNUPROC.SYSUT1 DD DSN=XMLR3.SYSUT1,
//          DISP=(MOD,DELETE,CATLG),
//          SPACE=(TRK,(5,5),RLSE),
//          UNIT=SYSDA
//DSNUPROC.SORTOUT DD DSN=XMLR3.SORTOUT,
//          DISP=(MOD,DELETE,CATLG),
//          SPACE=(TRK,(5,5),RLSE),
//          UNIT=SYSDA
//DSNUPROC.SYSIN DD *
REORG TABLESPACE XMLR3DB.TSAUDIT1
        STATISTICS TABLE(ALL)
        INDEX(ALL)
REORG TABLESPACE XMLR3DB.XBKRO000
        STATISTICS TABLE(ALL)
        INDEX(ALL)
        WORKDDN(SYSUT1)
/*

```

16 KB buffer pool

XML table spaces are always defined in 16K buffer pool. You must monitor the size of the 16K buffer pool, and ensure that it is appropriately sized and backup by real storage.

DSNZPARM settings

XMLVALA and XMLVALS are DSNZPARMs which control the amount of virtual storage in the XML pool which is used as working storage for document materialization, and XPath evaluation.

As described in 10.1.3, “Sizing XMLVALA and XMLVALS” on page 231, the XMLVALA subsystem parameter specifies, in KB, an upper limit for the amount of storage that each user is to have for storing XML values.

- ▶ Acceptable values: 1 to 2,097,152 KB
- ▶ Default: 204,800 KB

The XMLVALS subsystem parameter specifies, in MB, an upper limit for the amount of storage that each system can use for storing XML values.

- ▶ Acceptable values: 1 to 51200 MB
- ▶ Default: 10240 MB

Bear in mind that these values may need to be adjusted if your application materializes a large volume of XML data.

RID pool size

The RID pool size may need to be increased because XML index access (particularly DOCID ANDing or ORing) also uses RID pool.

11.7 SQL/XML coding techniques

Coding SQL/XML will be a new sport for many established DB2 programmers. With several new concepts to learn in the SQL/XML language, it is sometimes hard to work out “how” to write the SQL/XML statement that you want. And then you will want to work out what programming techniques will actually yield the best performance. This section provides a collection of some common coding techniques that will yield efficient data access.

11.7.1 XMLTABLE to minimize database calls

The XMLTABLE function is a very powerful way to minimize the number of DB2 database calls to access XML data. It is possible to retrieve XML data elements using the XMLQUERY function multiple times. However it is more efficient to replace multiple XMLQUERY calls with a single XMLTABLE call, as shown in Example 11-10.

Example 11-10 Multiple XMLQUERY calls replaced with a single XMLTABLE call

```

SET V_CREDITM = (
  select xmlcast(xmlquery('declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $d/Document/BkToCstmrStmt/GrpHdr/CreDtTm'
    passing VALIDXML as "d") as timestamp) from sysibm.sysdummy1 );

SET V_MINISTMT = (
  select xmlquery('declare default element namespace
    "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $d/Document/BkToCstmrStmt/Stmt'
    passing VALIDXML as "d") from SYSIBM.SYSDUMMY1 );

-- can be replaced by

SELECT X.CRE_DT_TM, X.MINISTMT INTO V_CREDITM, V_MINISTMT
  FROM XMLTable(XMLNAMESPACES(DEFAULT
    'urn:iso:std:iso:20022:tech:xsd:camt.053.001.02',
    '$d/Document/BkToCstmrStmt' PASSING VALIDXML as "d"
    COLUMNS
      "CRE_DT_TM" TIMESTAMP PATH './GrpHdr/CreDtTm/text()'
      "MINISTMT" XML PATH './Stmt' ) AS X ;

```

Another way of making the code more elegant and more efficient would be to combine the two XMLQUERY calls into a single select statement, as shown in Example 11-11.

Example 11-11 single select statement combining two xmlquery expressions

```

select
  xmlcast(xmlquery('
    declare default element namespace
      "urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
    $d/Document/BkToCstmrStmt/GrpHdr/CreDtTm'

```

```
    passing VALIDXML as "d") as timestamp)
      into V_CREDITM,
xmlquery('declare default element namespace
"urn:iso:std:iso:20022:tech:xsd:camt.053.001.02";
$d/Document/BkToCstmrStmt/Stmt'
    passing VALIDXML as "d")
      into V_MINISTMT
from sysibm.sysdummy1 );
```

11.7.2 XMLEXISTS for index access

It really does not matter whether you use an XML or a relational index in order to reduce the number of rows that need to be accessed in the base table. Use XMLEXISTS whenever possible to help the optimizer. If there is no suitable relational index, then using an XML index will require the XMLEXISTS function in the vast majority of cases.

Using predicates to filter rows is a good practice in any situation, but since XMLQUERY and XMLTABLE can be more CPU-heavy than other built-in functions, good filtering is critical.


11.7.3 Simple XPath expressions

As a generalization, "simple" XPath expressions perform much better than "complex" XPath expressions, because they are more likely to qualify for XML index access path selection.

For example, XPath with / will generally perform better than with //, both for queries and XML index specifications.

It is also well worth reviewing the latest APARs to make sure that you download PTFs that improve performance of XML processing. The best place to start is II14426, info APAR to link together all the XML support delivery APARs.

<http://www.ibm.com/support/docview.wss?uid=isg1II14426>



Application scenario documents

There are many international organizations that publish XML standards for various industries. We have chosen to use the Bank To Customer Statement V2, one of the ISO 20022 (Universal financial industry message scheme) as the openly published XML standard for the XML documents. Our scenario is based on receiving and processing one of those messages from a financial institution such as a bank. You can read all about the ISO 20022 Universal financial Industry message scheme at:

<http://www.iso20022.org>

A full description of the message can be found in "Payments_Maintenance_2009.pdf" document which can be downloaded from:

http://www.iso20022.org/documents/general/Payments_Maintenance_2009.zip.

A.1 Schema

We have used the Bank To Customer Statement V2 schema which can be downloaded from <http://www.iso20022.org/documents/messages/camt/schemas/camt.053.001.02.zip>.

This schema is also available as additional material download file as described in Appendix B, "Additional material" on page 273.

A.2 XML message

The XML message was taken from <http://www.iso20022.org/documents/messages/camt/instances/camt.053.001.02.zip> and was augmented by adding a second Stmt element just like the first one with a few text values changed. The updated XML message is available as additional materials download file as described in Appendix B, "Additional material" on page 273.

We have provided two copies of the message.

Example A-1 shows the entire message which is inserted into BK_TO_CUST_STMT table intact.

The other is a shorter version containing only the elements that we actually process.

Example A-1 XML message received

```
<?xml version="1.0" encoding="UTF-8" ?>
  <Document xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02">
    <BkToCstmrStmt>
      <GrpHdr>
        <MsgId>AAAASESS-FP-STAT001</MsgId>
        <CreDtTm>2010-10-18T17:00:00+01:00</CreDtTm>
        <MsgPgntn>
          <PgNb>1</PgNb>
          <LastPgInd>true</LastPgInd>
        </MsgPgntn>
      </GrpHdr>
      <Stmt>
        <Id>AAAASESS-FP-STAT001</Id>
        <CreDtTm>2010-10-18T17:00:00+01:00</CreDtTm>
        <FrToDt>
          <FrDtTm>2010-10-18T08:00:00+01:00</FrDtTm>
          <ToDtTm>2010-10-18T17:00:00+01:00</ToDtTm>
        </FrToDt>
        <Acct>
          <Id>
            <Othr>
              <Id>50000000054910000003</Id>
            </Othr>
          </Id>
          <Ownr>
            <Nm>FINPETROL</Nm>
          </Ownr>
          <Svcr>
            <FinInstnId>
              <Nm>AAAA BANKEN</Nm>
              <PstlAdr>
```



```

        <Ctry>SE</Ctry>
      </PstlAdr>
    </FinInstnId>
  </Svcr>
</Acct>
<Bal>
  <Tp>
    <CdOrPrtry>
      <Cd>OPBD</Cd>
    </CdOrPrtry>
  </Tp>
  <Amt Ccy="SEK">500000</Amt>
  <CdtDbtInd>CRDT</CdtDbtInd>
  <Dt>
    <Dt>2010-10-15</Dt>
  </Dt>
</Bal>
<Bal>
  <Tp>
    <CdOrPrtry>
      <Cd>CLBD</Cd>
    </CdOrPrtry>
  </Tp>
  <Amt Ccy="SEK">435678.50</Amt>
  <CdtDbtInd>CRDT</CdtDbtInd>
  <Dt>
    <Dt>2010-10-18</Dt>
  </Dt>
</Bal>
<Ntry>
  <Amt Ccy="SEK">105678.50</Amt>
  <CdtDbtInd>CRDT</CdtDbtInd>
  <Sts>BOOK</Sts>
  <BookgDt>
    <DtTm>2010-10-18T13:15:00+01:00</DtTm>
  </BookgDt>
  <ValDt>
    <Dt>2010-10-18</Dt>
  </ValDt>
  <AcctSvcrRef>AAAASESS-FP-CN_98765/01</AcctSvcrRef>
  <BkTxCd>
    <Domn>
      <Cd>PAYM</Cd>
      <Fmly>
        <Cd>0001</Cd>
        <SubFmlyCd>0005</SubFmlyCd>
      </Fmly>
    </Domn>
  </BkTxCd>
  <NtryDtls>
    <TxDtls>
      <Refs>
        <EndToEndId>MUELL/FINP/RA12345</EndToEndId>
      </Refs>
      <RltdPties>
        <Dbtr>
          <Nm>MUELLER</Nm>
        </Dbtr>
      </RltdPties>
    </TxDtls>
  </NtryDtls>
</Ntry>

```

```

    </NtryDt1s>
  </Ntry>
  <Ntry>
    <Amt Ccy="SEK">200000</Amt>
    <CdtDbtInd>DBIT</CdtDbtInd>
    <Sts>BOOK</Sts>
    <BookgDt>
      <DtTm>2010-10-18T10:15:00+01:00</DtTm>
    </BookgDt>
    <ValDt>
      <Dt>2010-10-18</Dt>
    </ValDt>
    <AcctSvcrRef>AAAASESS-FP-ACCR-01</AcctSvcrRef>
    <BkTxCd>
      <Domn>
        <Cd>PAYM</Cd>
        <Fmly>
          <Cd>0001</Cd>
          <SubFmlyCd>0003</SubFmlyCd>
        </Fmly>
      </Domn>
    </BkTxCd>
    <NtryDt1s>
      <Btch>
        <MsgId>FINP-0055</MsgId>
        <PmtInfId>FINP-0055/001</PmtInfId>
        <NbOfTxS>20</NbOfTxS>
      </Btch>
    </NtryDt1s>
  </Ntry>
  <Ntry>
    <Amt Ccy="SEK">30000</Amt>
    <CdtDbtInd>CRDT</CdtDbtInd>
    <Sts>BOOK</Sts>
    <BookgDt>
      <DtTm>2010-10-18T15:15:00+01:00</DtTm>
    </BookgDt>
    <ValDt>
      <Dt>2010-10-18</Dt>
    </ValDt>
    <AcctSvcrRef>AAAASESS-FP-CONF-FX</AcctSvcrRef>
    <BkTxCd>
      <Domn>
        <Cd>TREA</Cd>
        <Fmly>
          <Cd>0002</Cd>
          <SubFmlyCd>0000</SubFmlyCd>
        </Fmly>
      </Domn>
    </BkTxCd>
    <NtryDt1s>
      <TxDt1s>
        <Refs>
          <InstrId>FP-004567-FX</InstrId>
          <EndToEndId>AAAASS1085FINPSS</EndToEndId>
        </Refs>
        <AmtDt1s>
          <CntrValAmt>
            <Amt Ccy="EUR">3255</Amt>
            <CcyXchg>

```

```

        <SrcCcy>EUR</SrcCcy>
        <XchgRate>0.1085</XchgRate>
    </CcyXchg>
</CntrValAmt>
</AmtDt1s>
</TxDt1s>
</NtryDt1s>
</Ntry>
</Stmt>
<Stmt>
    <Id>AAAASESS-FP-STAT002</Id>
    <CreDtTm>2010-10-17T17:00:00+01:00</CreDtTm>
    <FrToDt>
        <FrDtTm>2010-10-17T08:00:00+01:00</FrDtTm>
        <ToDtTm>2010-10-17T17:00:00+01:00</ToDtTm>
    </FrToDt>
    <Acct>
        <Id>
            <Othr>
                <Id>50000000054910000004</Id>
            </Othr>
        </Id>
        <Ownc>
            <Nm>FINPETROL</Nm>
        </Ownc>
        <Svcr>
            <FinInstnId>
                <Nm>AAAB BANKEN</Nm>
                <Pst1Adr>
                    <Ctry>SE</Ctry>
                </Pst1Adr>
            </FinInstnId>
        </Svcr>
    </Acct>
    <Bal>
        <Tp>
            <CdOrPrtry>
                <Cd>OPAV</Cd>
            </CdOrPrtry>
        </Tp>
        <Amt Ccy="SEK">500300</Amt>
        <CdtDbtInd>CRDT</CdtDbtInd>
        <Dt>
            <Dt>2010-10-14</Dt>
        </Dt>
    </Bal>
    <Bal>
        <Tp>
            <CdOrPrtry>
                <Cd>FWAV</Cd>
            </CdOrPrtry>
        </Tp>
        <Amt Ccy="SEK">435478.50</Amt>
        <CdtDbtInd>CRDT</CdtDbtInd>
        <Dt>
            <Dt>2010-10-17</Dt>
        </Dt>
    </Bal>
    <Ntry>
        <Amt Ccy="SEK">105378.50</Amt>

```

```

<CdtDbtInd>CRDT</CdtDbtInd>
<Sts>BOOK</Sts>
<BookgDt>
  <DtTm>2010-10-17T13:15:00+01:00</DtTm>
</BookgDt>
<ValDt>
  <Dt>2010-10-17</Dt>
</ValDt>
<AcctSvcrRef>AAAASESS-FP-CN_98764/01</AcctSvcrRef>
<BkTxCd>
  <Domn>
    <Cd>PAYM</Cd>
    <Fmly>
      <Cd>0002</Cd>
      <SubFmlyCd>0004</SubFmlyCd>
    </Fmly>
  </Domn>
</BkTxCd>
<NtryDtls>
  <TxDtls>
    <Refs>
      <EndToEndId>MUELL/FINP/RA12344</EndToEndId>
    </Refs>
    <RltdPties>
      <Dbtr>
        <Nm>MUELLAR</Nm>
      </Dbtr>
    </RltdPties>
  </TxDtls>
</NtryDtls>
</Ntry>
<Ntry>
  <Amt Ccy="SEK">200100</Amt>
  <CdtDbtInd>DBIT</CdtDbtInd>
  <Sts>BOOK</Sts>
  <BookgDt>
    <DtTm>2010-10-17T10:15:00+01:00</DtTm>
  </BookgDt>
  <ValDt>
    <Dt>2010-10-17</Dt>
  </ValDt>
  <AcctSvcrRef>AAAASESS-FP-ACCR-02</AcctSvcrRef>
  <BkTxCd>
    <Domn>
      <Cd>TREA</Cd>
      <Fmly>
        <Cd>0002</Cd>
        <SubFmlyCd>0004</SubFmlyCd>
      </Fmly>
    </Domn>
  </BkTxCd>
  <NtryDtls>
    <Btch>
      <MsgId>FINP-0056</MsgId>
      <PmtInfId>FINP-0055/002</PmtInfId>
      <NbOfTx>21</NbOfTx>
    </Btch>
  </NtryDtls>
</Ntry>
<Ntry>

```

```

<Amt Ccy="SEK">30020</Amt>
<CdtDbtInd>CRDT</CdtDbtInd>
<Sts>BOOK</Sts>
<BookgDt>
  <DtTm>2010-10-17T15:15:00+01:00</DtTm>
</BookgDt>
<ValDt>
  <Dt>2010-10-17</Dt>
</ValDt>
<AcctSvcrRef>AAAASESS-FP-CONF-FY</AcctSvcrRef>
<BkTxCd>
  <Domn>
    <Cd>TREA</Cd>
    <Fmly>
      <Cd>0003</Cd>
      <SubFmlyCd>0001</SubFmlyCd>
    </Fmly>
  </Domn>
</BkTxCd>
<NtryDtls>
  <TxDtls>
    <Refs>
      <InstrId>FP-004568-FX</InstrId>
      <EndToEndId>AAAASS1084FINPSS</EndToEndId>
    </Refs>
    <AmtDtls>
      <CntrValAmt>
        <Amt Ccy="EUR">3254</Amt>
        <CcyXchg>
          <SrcCcy>EUR</SrcCcy>
          <XchgRate>0.1084</XchgRate>
        </CcyXchg>
      </CntrValAmt>
    </AmtDtls>
  </TxDtls>
</NtryDtls>
</Ntry>
</Stmt>
</BkToCstmrStmt>
</Document>

```

Example A-2 is a shorter version of the Message Received and contains only the elements that we actually process in our scenario.

Example A-2 XML message parts processed

```

<?xml version="1.0" encoding="UTF-8" ?>
  <Document xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="urn:iso:std:iso:20022:tech:xsd:camt.053.001.02">
    <BkToCstmrStmt>
      <GrpHdr>
        <MsgId>AAAASESS-FP-STAT001</MsgId>
        <CreDtTm>2010-10-18T17:00:00+01:00</CreDtTm>
      </GrpHdr>
      <Stmt>
        <Id>AAAASESS-FP-STAT001</Id>
        <CreDtTm>2010-10-18T17:00:00+01:00</CreDtTm>
        <FrToDt>
          <FrDtTm>2010-10-18T08:00:00+01:00</FrDtTm>
          <ToDtTm>2010-10-18T17:00:00+01:00</ToDtTm>
        </FrToDt>
      </Stmt>
    </BkToCstmrStmt>
  </Document>

```

```

</FrToDt>
<Acct>
  <Id>
    <Othr>
      <Id>50000000054910000003</Id>
    </Othr>
  </Id>
  <Ownc>
    <Nm>FINPETROL</Nm>
  </Ownc>
  <Svcr>
    <FinInstnId>
      <Nm>AAAA BANKEN</Nm>
    </FinInstnId>
  </Svcr>
</Acct>
<Bal>
  <Tp>
    <CdOrPrtry>
      <Cd>OPBD</Cd>
    </CdOrPrtry>
  </Tp>
  <Amt Ccy="SEK">500000</Amt>
  <CdtDbtInd>CRDT</CdtDbtInd>
</Bal>
<Bal>
  <Tp>
    <CdOrPrtry>
      <Cd>CLBD</Cd>
    </CdOrPrtry>
  </Tp>
  <Amt Ccy="SEK">435678.50</Amt>
  <CdtDbtInd>CRDT</CdtDbtInd>
</Bal>
<Ntry>
  <Amt Ccy="SEK">105678.50</Amt>
  <BookgDt>
    <DtTm>2010-10-18T13:15:00+01:00</DtTm>
  </BookgDt>
  <AcctSvcrRef>AAAASESS-FP-CN_98765/01</AcctSvcrRef>
</Ntry>
<Ntry>
  <Amt Ccy="SEK">200000</Amt>
  <BookgDt>
    <DtTm>2010-10-18T10:15:00+01:00</DtTm>
  </BookgDt>
  <AcctSvcrRef>AAAASESS-FP-ACCR-01</AcctSvcrRef>
</Ntry>
<Ntry>
  <Amt Ccy="SEK">30000</Amt>
  <BookgDt>
    <DtTm>2010-10-18T15:15:00+01:00</DtTm>
  </BookgDt>
  <AcctSvcrRef>AAAASESS-FP-CONF-FX</AcctSvcrRef>
</Ntry>
</Stmt>
<Stmt>
  <Id>AAAASESS-FP-STAT002</Id>
  <CreDtTm>2010-10-17T17:00:00+01:00</CreDtTm>
  <FrToDt>

```

```

    <FrDtTm>2010-10-17T08:00:00+01:00</FrDtTm>
    <ToDtTm>2010-10-17T17:00:00+01:00</ToDtTm>
  </FrToDt>
  <Acct>
    <Id>
      <Othr>
        <Id>50000000054910000004</Id>
      </Othr>
    </Id>
    <Ownr>
      <Nm>FINPETROL</Nm>
    </Ownr>
    <Svcr>
      <FinInstnId>
        <Nm>AAAB BANKEN</Nm>
      </FinInstnId>
    </Svcr>
  </Acct>
  <Bal>
    <Tp>
      <CdOrPrtry>
        <Cd>OPAV</Cd>
      </CdOrPrtry>
    </Tp>
    <Amt Ccy="SEK">500300</Amt>
    <CdtDbtInd>CRDT</CdtDbtInd>
  </Bal>
  <Bal>
    <Tp>
      <CdOrPrtry>
        <Cd>FWAV</Cd>
      </CdOrPrtry>
    </Tp>
    <Amt Ccy="SEK">435478.50</Amt>
    <CdtDbtInd>CRDT</CdtDbtInd>
  </Bal>
  <Ntry>
    <Amt Ccy="SEK">105378.50</Amt>
    <CdtDbtInd>CRDT</CdtDbtInd>
    <ValDt>
      <Dt>2010-10-17</Dt>
    </ValDt>
    <AcctSvcrRef>AAAASESS-FP-CN_98764/01</AcctSvcrRef>
  </Ntry>
  <Ntry>
    <Amt Ccy="SEK">200100</Amt>
    <ValDt>
      <Dt>2010-10-17</Dt>
    </ValDt>
    <AcctSvcrRef>AAAASESS-FP-ACCR-02</AcctSvcrRef>
  </Ntry>
  <Ntry>
    <Amt Ccy="SEK">30020</Amt>
    <BookgDt>
      <DtTm>2010-10-17T15:15:00+01:00</DtTm>
    </BookgDt>
    <AcctSvcrRef>AAAASESS-FP-CONF-FY</AcctSvcrRef>
  </Ntry>
</Stmnt>
</BkToCstmrStmnt>

```

</Document>

**B**

Additional material

This book refers to additional material that can be downloaded from the Internet as described in the following sections.

Locating the Web material

The Web material associated with this book is available in softcopy on the Internet from the IBM Redbooks Web server. Point your Web browser at:

<ftp://www.redbooks.ibm.com/redbooks/SG247915>

Alternatively, you can go to the IBM Redbooks Web site at:

ibm.com/redbooks

Select the **Additional materials** and open the directory that corresponds with the IBM Redbooks publication number SG24-7915-00.

Using the Web material

The additional Web material that accompanies this book includes the following files:

<i>File name</i>	<i>Description</i>
Storedproceduresamples.zip	Zipped code for the setup of stored procedures used for the implementation of the scenarios
Javasamples.zip	Zipped code samples and data used for the implementation of the scenario in Java
Cobolsamples.zip	Zipped code samples and data used for the implementation of the scenario in COBOL

System requirements for downloading the Web material

The Web material requires the following system configuration:

Hard disk space:	100 KB
Operating System:	Windows
Processor:	All Intel® and AMD processors capable of running the supported Windows operating systems (32-bit and x64 based systems).
Memory:	2 GB

Downloading and extracting the Web material

Create a subdirectory (folder) on your workstation, and extract the contents of the Web material .zip file into this folder.

Storedproceduresamples.zip

The zipped file contain code and data to reproduce the application infrastructure.

Prerequisites

- ▶ DB2 Client for Windows V9.7 Fix Pack 3a
- ▶ IBM Data Server Driver for JDBC and SQLJ version 4.9 or later (If you use DB2 Client, V9.7 Fix Pack 3a and above provides this support)
- ▶ Environment setup on server for XML schema registration
- ▶ The table BK_TO_CSTMR_STMT with XML column, and other related tables documented in “Tables used for following examples.” on page 88.

Sample code

- ▶ STOREXML1.db2
Stored procedure in Example 6-3 on page 89.
- ▶ Teststorexml1.java
Java program to drive STOREXML1().
- ▶ callstorexml1.bat
Script to run Teststorexml1 java program.
- ▶ STOREXML2.db2
Stored procedure in Example 6-4 on page 91.
- ▶ Teststorexml2.java
Java program to drive STOREXML2().
- ▶ callstorexml2.bat
Script to run Teststorexml2 Java program.
- ▶ LOADMQ.db2
Stored procedure to prime an MQ queue with an XML message.
- ▶ STOREXML3.db2
Stored procedure in Example 6-7 on page 96.
- ▶ STOREXML4.db2
Stored procedure in Example 6-11 on page 98.
- ▶ STOREXML5.db2
Stored procedure in Example 6-25 on page 114.
- ▶ Teststorexml5.java

- ▶ Java program to drive STOREXML5().
- ▶ callstorexml5.bat
 - ▶ Script to run Teststorexml5 java program
- ▶ RECEIVE_CDC.db2
 - ▶ Stored procedure in Example 6-34 on page 122.
- ▶ Testcdc.java
 - ▶ Java program to drive RECEIVE_CDC().
- ▶ callcdc.bat
 - ▶ Script to run Testcdc Java program.

Sample data

- ▶ XMLTABLES.DDL
 - ▶ DDL for all samples.
- ▶ XML_TEST_SOURCE.SQL
 - ▶ Test data for all samples.
- ▶ RESET_TEST_SOURCE.SQL
 - ▶ Script to reset test data for all samples.

Javasamples.zip

The zipped file contain code and data to reproduce our Java scenario.

Prerequisites

- ▶ SDK for Java Version 6 or later
- ▶ IBM Data Server Driver for JDBC and SQLJ version 4.9 or later (If you use DB2 Client, V9.7 Fix Pack 3a and above provides this support)
- ▶ Environment setup on server for XML schema registration
- ▶ The table BK_TO_CSTMR_STMT with XML column should be created as Example 7-5 on page 139.

Sample code

- ▶ RegisterSchema.java
 - ▶ This class registers the XML schema to the DB2 databases.
- ▶ InsertXML.java
 - ▶ This class validates and inserts XML data into db2 table.
- ▶ UpdateXML.java
 - ▶ This class demonstrates partial updates of XML documents.
- ▶ SelectXML.java
 - ▶ This class demonstrates the retrieving entire or partial XML document to a SQLXML object.
- ▶ XMLProcedure.java
 - ▶ This class creates SQL stored procedure with XML as parameter to shred XML document, then calls the SQL stored procedure from Java.
- ▶ TransformXML.java

This class transforms retrieved XML document into an new XML or HTML document.

- ▶ SendMQMessage.java

This class sends a XML message generated by XSLT to WebSphere MQ.

Sample data

- ▶ Schema of bank to customer statement message
camt.053.001.02.xsd
- ▶ XML file of bank to customer statement message
camt.053.001.02.xml
camt.053.001.03.xml
camt.053.001.04.xml
- ▶ XSLT file to transform XML message
camt.053.001.04.xsl
camt.053.001.05.xsl
- ▶ XML document send to WebSphere MQ
xmloutput.xml

Cobolsamples.zip

The zipped file contain code and data to reproduce our COBOL scenario.

Prerequisites

- ▶ Enterprise COBOL for z/OS Version 4.1 or later.
- ▶ DB2 Connect Version 9.5 or later.
- ▶ Configuration of DB2 for z/OS node in DB2 Connect.
- ▶ Setup of DB2 XML schema repository.
- ▶ The XML schema should be registered as in Example 8-8 on page 164.
- ▶ The table BK_TO_CSTMR_STMT with XML column should be created as Example 8-9 on page 164.

Sample code

- ▶ INSBKST.txt
COBOL program for inserting a BankToCustomerStatement into table BK_TO_CUSTMR_STSMT using a file reference variable.
- ▶ GETSTMT.txt
COBOL program for selecting a BankToCustomerStatement from table BK_TO_CUSTMR_STSMT into a file using a file reference variable.
- ▶ GETNTRY.txt
COBOL program for selecting entries from a BankToCustomerStatement from table BK_TO_CUSTMR_STSMT using XMLTABLE.
- ▶ UPDRCP2.txt
COBOL program for altering the message recipient of a BankToCustomerStatement in table BK_TO_CUSTMR_STSMT using XMLMODIFY. The message recipient element is received as a file reference variable.
- ▶ UPDRCP2.txt

COBOL program for altering the message recipient of a BankToCustomerStatement in table BK_TO_CUSTMR_STSMT using XMLMODIFY. The message recipient element is generated using COBOL XML GENERATE statement.

▶ PREPPROG.txt

JCL for precompiling, compiling, link-editing and binding COBOL program INSBKST. It can be modified to prepare other COBOL programs by replacing INSBKST with other COBOL program name.

▶ RUNPROG.txt

JCL for running COBOL program INSBKST. It can be modified to run other COBOL programs by replacing the program name, but take care to enter the correct input variables.

Sample data

▶ Schema for BankToCustomerStatement message

camt.053.001.02.xsd

▶ XML file with BankToCustomerStatement message

camt.stmt.xml

▶ XML file with MsgRcpt element

camt.msgrcpt.xml

Glossary

A

address space. A range of virtual storage pages identified by a number (ASID) and a collection of segment and page tables which map the virtual pages to real pages of the computer's memory.

address space connection. The result of connecting an allied address space to DB2. Each address space containing a task connected to DB2 has exactly one address space connection, even though more than one task control block (TCB) can be present. See allied address space and task control block.

Advanced Program-to-Program communication (APPC). (1) The general facility characterizing the LU6.2 architecture and its implementation in different SNA products. (2) Sometimes used to refer to an LU6.2 product feature in particular, such as an APPC application programming interface.

allied address space. An area of storage external to DB2 that is connected to DB2 and is therefore capable of requesting DB2 services.

American National Standards Institute (ANSI). An organization consisting of producers, consumers, and general interest groups, that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States.

ANSI. See *American National Standards Institute*.

APAR. See *authorized program analysis report*.

API. See *Application Program Interface*.

applet. See *Java Applet*.

application. (1) A program or set of programs that perform a task; for example, a payroll application. (2) In Java programming, a self-contained, stand-alone Java program that includes a static main method. It does not require an applet viewer. Contrast with *applet*.

application plan. The control structure produced during the bind process and used by DB2 to process SQL statements encountered during statement execution.

application program interface (API). A functional interface supplied by the operating system or by a separately orderable licensed program that allows an application program written in a high-level language to use specific data or functions of the operating system or licensed program.

application requester (AR). See *requester*.

Application Service Provider (ASP). An ASP is an agent or broker that aggregates, facilitates and brokers IT services to deliver IT-enabled business solutions across a network via subscription-based pricing.

application-owning region (AOR). A CICS® region in an MRO environment that "owns" the CICS applications, and invokes them on behalf of remotely attached terminal (or Web) users. See also *TOR* and *listener region*.

AR. Application requester. See *requester*.

ASCII. (1) American Standard Code for Information Interchange. A standard assignment of 7-bit numeric codes to characters. See also *Unicode*. (2) An encoding scheme used to represent strings in many environments, typically on PCs and workstations. Contrast with *EBCDIC*.

attribute. In XML, a name="value" pair that can be placed in the start tag of an element. The value must be quoted with single or double quotes.

authorization ID. A string that can be verified for connection to DB2 and to which a set of privileges are allowed. It can represent an individual, an organizational group, or a function, but DB2 does not determine this representation.

authorized program analysis report (APAR). A report of a problem caused by a suspected defect in a current, unaltered release of a program.

automatic bind. (More correctly automatic rebind). A process by which SQL statements are bound automatically (without a user issuing a BIND command) when an application process begins execution and the bound application plan or package it requires is not valid.

B

base table. (1) A table created by the SQL CREATE TABLE statement that is used to hold persistent data. Contrast with *result table* and *temporary table*. (2) A table containing a LOB column definition. The actual LOB column data is not stored along with the base table. The base table contains a row identifier for each row and an indicator column for each of its LOB columns. Contrast with *auxiliary table*.

basic mode. A S/390® central processing mode that does not use logical partitioning. Contrast with *logically partitioned (LPAR) mode*.

bean. A definition or instance of a JavaBeans component. See *JavaBeans*.

binary XML format. A system of storing XML data in binary, as opposed to text, that facilitates more efficient storage and exchange.

bind. The process by which the output from the DB2 precompiler is converted to a usable control structure called a package or an application plan. During the process, access paths to the data are selected and some authorization checking is performed.

browser. (1) In VisualAge® for Java, a window that provides information on program elements. There are browsers for projects, packages, classes, methods, and interfaces. (2) An Internet-based tool that lets users browse Web sites.

bytecode. Machine-independent code generated by the Java compiler and executed by the Java interpreter.

C

call level interface (CLI). A callable application program interface (API) for database access, which is an alternative to using embedded SQL. In contrast to embedded SQL, DB2 CLI does not require the user to precompile or bind applications, but instead provides a standard set of functions to process SQL statements and related services at run time.

Cascading Style Sheet (CSS). CSS defines a stylesheet language for HTML 4.0. CSS allows a Web page designer to separately specify style elements of a Web page, such as colors, fonts and font styles.

case-sensitive. Indicates whether an application, processor, or operating system distinguishes between upper and lower case. If it does, it is case-sensitive. XML tags are case-sensitive, but HTML tags are not.

casting. Explicitly converting an object or primitive's data type.

catalog. In DB2, a collection of tables that contains descriptions of objects such as tables, views, and indexes.

catalog table. Any table in the DB2 catalog.

CGI. The Common Gateway Interface (CGI) is a means of allowing a Web server to execute a program that you provide rather than to retrieve a file. A number of popular Web servers support the CGI. For some applications (for example, displaying information from a database), you must do more than simply retrieve an HTML document from a disk and send it to the Web browser. For such applications, the Web server has to call a program to generate the HTML to be displayed. The CGI is not the only such interface, however.

channel-attached. (1) Pertaining to attachment of devices directly by data channels (I/O channels) to a computer. (2) Pertaining to devices attached to a controlling unit by cables rather than by telecommunication lines.

character large object (CLOB). See CLOB.

class. An encapsulated collection of data and methods to operate on the data. A class may be instantiated to produce an object that is an instance of the class.

class hierarchy. The relationships between classes that share a single inheritance. All Java classes inherit from the Object class.

class method. Methods that apply to the class as a whole rather than its instances (also called a *static method*).

class variable. Variables that apply to the class as a whole rather than its instances (also called a *static field*).

CLASSPATH. In your deployment environment, the environment variable keyword that specifies the directories in which to look for class and resource files.

CLI. See *call level interface*.

client. (1) A networked computer in which the IDE is connected to a repository on a team server. (2) See requester.

CLOB. A sequence of bytes representing single-byte characters or a mixture of single and double-byte characters where the size can be up to 2 GB - 1. Although the size of character large object values can be anywhere up to 2 GB - 1, in general, they are used whenever a character string might exceed the limits of the VARCHAR type.

codebase. An attribute of the <APPLET> tag that provides the relative path name for the classes. Use this attribute when your class files reside in a different directory than your HTML files.

column function. An SQL operation that derives its result from a collection of values across one or more rows. Contrast with scalar function.

commit. The operation that ends a unit of work by releasing locks so that the database changes made by that unit of work can be perceived by other processes.

Common Connector Framework. In the Enterprise Access Builder, interface and class definitions that provide a consistent means of interacting with enterprise resources (for example, CICS and Encina transactions) from any Java execution environment.

connection. In the VisualAge for Java Visual Composition Editor, a visual link between two components that represents the relationship between the components. Each connection has a source, a target, and other properties.

connection handle. The data object that contains information associated with a connection managed by DB2 CLI. This includes general status information, transaction status, and diagnostic information.

content model. In XML, the expression specifying what elements and data are allowed within an element.

cookie. (1) A small file stored on an individual's computer; this file allows a site to tag the browser with a unique identification. When a person visits a site, the site's server requests a unique ID from the person's browser. If this browser does not have an ID, the server delivers one. On the Wintel platform, the cookie is delivered to a file called 'cookies.txt,' and on a Macintosh platform, it is delivered to 'MagicCookie.' Just as someone can track the origin of a phone call with Caller ID, companies can use cookies to track information about behavior. (2) Persistent data stored by the client in the Servlet Builder.

cursor. A named control structure used by an application program to point to a row of interest within some set of rows, and to retrieve rows from the set, possibly making updates or deletions.

Customer relationship management (CRM). CRM includes the systems and infrastructure required to analyze, capture and share all parts of the customer's relationship with the enterprise. From a strategy perspective, it represents a process to measure and allocate organizational resources to those activities that have the greatest return and impact on profitable customer relationships.

D

Data Access Bean. In the VisualAge for Java Visual Composition Editor, a bean that accesses and manipulates the content of JDBC/ODBC-compliant relational databases.

Data Access Builder. A VisualAge for Java Enterprise tool that generates beans to access and manipulate the content of JDBC/ODBC-compliant relational databases.

data source. A local or remote relational or non-relational data manager that is capable of supporting data access via an ODBC driver which supports the ODBC APIs. In the case of DB2 for OS/390®, the data sources are always relational database managers.

database management system (DBMS). A software system that controls the creation, organization, and modification of a database and access to the data stored within it.

DB2 thread. The DB2 structure that describes an application's connection, traces its progress, processes resource functions, and delimits its accessibility to DB2 resources and services.

DBCLOB. A sequence of bytes representing double-byte characters where the size can be up to 2 gigabytes. Although the size of double-byte character large object values can be anywhere up to 2 gigabytes, in general, they are used whenever a double-byte character string might exceed the limits of the VARGRAPHIC type.

DBMS. Database management system.

direct access storage device (DASD). A mass storage medium on which a computer stores data.

distributed relational database architecture (DRDA)). A connection protocol for distributed relational database processing that is used by IBM's relational database products. DRDA includes protocols for communication between an application and a remote relational database management system, and for communication between relational database management systems.

DLL (dynamic link library). A file containing executable code and data bound to a program at load time or run time, rather than during linking. The code and data in a dynamic link library can be shared by several applications simultaneously. The DLL's Enterprise Access Builders also generate platform-specific DLLs for the workstation and OS/390 platforms.

Document Object Model (DOM). This allows the representation and manipulation of an XML document in memory as a programming object. DOM is defined by the World-Wide Web Consortium.

Document Type Definition (DTD). A DTD is a definition of which Elements and Attributes are acceptable in a specific XML file. The DTD therefore defines a subset of XML which may be used for a particular application.

DOM. (see Document Object Model).

DOM Tree. A DOM Tree is an in-memory representation of an XML Document.

double precision. A floating-point number that contains 64 bits. See also *single precision*.

double-byte character large object (DBCLOB). See DBCLOB.

DRDA. Distributed relational database architecture.

duplex. Pertaining to communication in which data or control information can be sent and received at the same time. Contrast with *half duplex*.

dynamic bind. A process by which SQL statements are bound as they are entered.

Dynamic I/O Reconfiguration. A S/390 function that allows I/O configuration changes to be made non-disruptively to the current operating I/O configuration.

dynamic SQL. SQL statements that are prepared and executed within an application program while the program is executing. In dynamic SQL, the SQL source is contained in host language variables rather than being coded into the application program. The SQL statement can change several times during the application program's execution.

E

EBCDIC. Extended binary coded decimal interchange code. An encoding scheme used to represent character data in the MVS™, VM, VSE, and OS/400® environments. Contrast with ASCII.

EBNF. Extended Backus-Naur Form. A formal set of production rules that comprise a grammar defining another language, such as XML.

Electronic data interchange. The automatic machine-to-machine transfer of trading documents (for example, invoices and purchase orders) using electronic networks such as the Internet. Originally conducted only through value-added networks, EDI is gradually moving to the Internet.

element. In XML, a start tag and its end tag, plus the content between the tags. An empty tag is also an element.

embedded SQL. SQL statements coded within an application program. See static SQL.

embeddedJava. An API and application environment for high-volume embedded devices, such as mobile phones, pagers, process control, instrumentation, office peripherals, network routers and network switches. EmbeddedJava applications run on real-time operating systems and are optimized for the constraints of small-memory footprints and diverse visual displays.

empty declaration. In XML, the DTD declaration for an empty tag. For example, if <foo/> is an empty tag, the empty declaration looks like: <!ELEMENT foo EMPTY>.

empty tag. In XML, a start and end tag combined in one tag. The tag has a trailing slash, so an XML parser can immediately recognize it as an empty tag and not bother looking for a matching end tag. For example, if foo is an empty tag, it looks like <foo/>.

Enterprise Java. Includes Enterprise JavaBeans as well as open API specifications for: database connectivity, naming and directory services, CORBA/IIOP interoperability, pure Java distributed computing, messaging services, managing system and network resources, and transaction services.

Enterprise JavaBeans. A cross-platform component architecture for the development and deployment of multi-tier, distributed, scalable, object-oriented Java applications.

Enterprise JavaBeans (EJB). The Enterprise JavaBeans specification defines a way of building transactionally aware business objects in Java.

Enterprise Systems Architecture/390 (ESA/390). An IBM architecture for mainframe computers and peripherals. Processors that follow this architecture include the S/390 Server family of processors.

entity. In XML, an entity declaration provides the ability to have constants or replacement strings, which are expanded by a pre-processor. An entity declaration maps some token to a replacement string. Later the token can be prefixed with the '&' character and the replacement string is put in its place.

environment handle. In DB2 ODBC, the data object that contains global information regarding the state of the application. An environment handle must be allocated before a connection handle can be allocated. Only one environment handle can be allocated per application.

ESA/390. See *Enterprise Systems Architecture/390*.

exception. An exception is an object that has caused some sort of new condition, such as an error. In Java, *throwing* an exception means passing that object to an interested party; a signal indicates what kind of condition has taken place. *Catching* an exception means receiving the sent object. *Handling* this exception usually means taking care of the problem after receiving the object, although it might mean doing nothing (which would be bad programming practice).

executable content. Code that runs from within an HTML file (such as an applet).

extends. A subclass or interface extends a class or interface if it add fields or methods, or overrides its methods.

external function. A function for which the body is written in a programming language that takes scalar argument values and produces a scalar result for each invocation. Contrast with sourced function and built-in function.

Extranet. In some cases intranets have connections to other independent intranets. An example would be one company connecting its intranet to the intranet of one of its suppliers. Such a connection of intranets is called an extranet. Depending on the implementation, they may or may not be fully or partially visible to the outside.

F

factory. A bean that dynamically creates instances of beans.

FastCGI. FastCGI is a way of combining the advantages of CGI programming with some of the performance benefits you get by using the GWAPI. FastCGI, written by Open Market, Inc., is an extension to normal Web server processing. It requires server-specific API support, which is available for AIX®, Sun Solaris, HP-UX, and OS/390. With FastCGI you can start applications in independent address spaces and pass requests for these applications from the Web server. The communication is through either the TCP/IP sockets interface or UNIX Domain socket bind path in the Hierarchical File System (HFS).

fibre channel standard. An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. The protocol has four layers. The lower of the four layers defines the physical media and interface, the upper of the four layers defines one or more logical protocols (for example, FCP for SCSI command protocols and FC-SB-2 for FICON® for ESA/390). Refer to ANSI X3.230.1999x.

FICON. (1) An ESA/390 computer peripheral interface. The I/O interface uses ESA/390 logical protocols over a FICON serial interface that configures attached units to a FICON communication fabric. (2) An FC4 proposed standard that defines an effective mechanism for the export of the SBCON command protocol via fibre channels.

field. A data object in a class; for example, a variable.

File Transfer Protocol (FTP). In the Internet suite of protocols, an application layer protocol that uses TCP and Telnet services to transfer bulk-data files between machines or hosts.

first tier. The client; the hardware and software with which the end user interacts.

foreign key. A key that is specified in the definition of a referential constraint. Because of the foreign key, the table is a dependent table. The key must have the same number of columns, with the same descriptions, as the primary key of the parent table.

form data. A generated class representing the HTML form elements in a visual servlet.

FTP. See *File Transfer Protocol*.

function. A specific purpose of an entity or its characteristic action, such as a column function or scalar function. (See column function and scalar function.) Furthermore, functions can be user-defined, built-in, or generated by DB2. (See user-defined function, external function, sourced function.)

G

garbage collection. Java's ability to clean up inaccessible unused memory areas ("garbage") on the fly. Garbage collection slows performance, but keeps the machine from running out of memory.

GWAPI. Because CGI has some architectural limitations, most Web servers provide an equivalent mechanism that is optimized for their native environment. Domino® Go Web Server, IBM's strategic Web server, offers the Domino Go Web Server Application Programming Interface (GWAPI), optimized for a given environment, such as OS/390. The GWAPI enables you to create dynamic content similar to the CGI, but in a more specialized way than the generalized CGI. The GWAPI process is similar to OS/390 exit processing. There is an exit point for various server functions that can be exploited.

H

half duplex. In data communication, pertaining to transmission in only one direction at a time. Contrast with *duplex*.

handle. In DB2 CLI, a variable that refers to a data structure and associated resources. See connection handle, environment handle.

hard disk drive. (1) A storage media within a storage server used to maintain information that the storage server requires. (2) A mass storage medium for computers that is typically available as a fixed disk or a removable cartridge.

hierarchy. The order of inheritance in object-oriented languages. Each class in the hierarchy inherits attributes and behavior from its superclass, except for the top-level Object class.

HTTPS. HTTPS is a de facto standard developed by Netscape for making HTTP flows secure. Technically, it is the use of HTTP over SSL.

Hypertext Markup Language (HTML). A file format, based on SGML, for hypertext documents on the Internet. Allows for the embedding of images, sounds, video streams, form fields and simple text formatting. References to other objects are embedded using URLs, enabling readers to jump directly to the referenced document.

Hypertext Transfer Protocol (HTTP). The Internet protocol, based on TCP/IP, used to fetch hypertext objects from remote hosts.

I

IDE. See *Integrated Development Environment*.

Identifier. A unique name or address that identifies things such as programs, devices or systems.

initial program load (IPL). (1) The initialization procedure that causes an operating system to commence operation. (2) The process by which a configuration image is loaded into storage at the beginning of a work day or after a system malfunction. (3) The process of loading system programs and preparing a system to run jobs.

Integrated Development Environment (IDE). In VisualAge for Java, the set of windows that provide the user with access to development tools. The primary windows are the Workbench, Log, Console, Debugger, and Repository Explorer.

Internet. The vast collection of interconnected networks that use TCP/IP and evolved from the ARPANET of the late 1960s and early 1970s. The number of independent networks connected into this vast global net is growing daily.

Internet Protocol (IP). In the Internet suite of protocols, a connectionless protocol that routes data through a network or interconnected networks. IP acts as an intermediary between the higher protocol layers and the physical network. However, this protocol does not provide error recovery and flow control, and does not guarantee the reliability of the physical network.

interpreter. A tool that translates and executes code line-by-line.

Intranet. A private network inside a company or organization that uses the same kinds of software that you would find on the Internet, but that are only for internal use. As the Internet has become more popular, many of the tools used on the Internet are being used in private networks; for example, many companies have Web servers that are available only to employees.

IP. See *Internet Protocol*.

IPL. See *initial program load*.

J

JAR file format. JAR (Java Archive) is a platform-independent file format that aggregates many files into one. Multiple Java applets and their requisite components (.class files, images, sounds and other resource files) can be bundled in a JAR file and subsequently downloaded to a browser in a single HTTP transaction.

Java. An object-oriented programming language for portable, interpretive code that supports interaction among remote objects. Java was developed and specified by Sun Microsystems, Incorporated. The Java environment consists of the JavaOS, the Virtual Machines for various platforms, the object-oriented Java programming language, and several class libraries.

Java applet. A small Java program designed to run within a Web browser. It is downloadable and executable by a browser or network computer.

Java beans. Java's component architecture, developed by Sun, IBM, and others. The components, called Java beans, can be parts of Java programs, or they can exist as self-contained applications. Java beans can be assembled to create complex applications, and they can run within other component architectures (such as ActiveX and OpenDoc).

Java Development Kit (JDK). The set of Java technologies made available to licensed developers by Sun Microsystems. Each release of the JDK contains the following: the Java Compiler, Java Virtual Machine, Java Class Libraries, Java Applet Viewer, Java Debugger, and other tools.

Java Naming and Directory Interface (JNDI). A set of APIs that assist with the interfacing to multiple naming and directory services. (Definition copyright 1996-1999 Sun Microsystems, Inc. All Rights Reserved. Used by permission.)

Java Native Interface (JNI). A native programming interface that allows Java code running inside a Java Virtual Machine (VM) to interoperate with applications and libraries written in other programming languages, such as C and C++.

Java Platform. The Java Virtual Machine and the Java Core classes make up the Java Platform. The Java Platform provides a uniform programming interface to a 100% Pure Java program regardless of the underlying operating system. (Definition copyright 1996-1999 Sun Microsystems, Inc. All Rights Reserved. Used by permission.)

Java Remote Method Invocation (RMI). Java Remote Method Invocation is method invocation between peers, or between client and server, when applications at both ends of the invocation are written in Java. Included in JDK 1.1.

Java Runtime Environment (JRE). A subset of the Java Development Kit for end users and developers who want to redistribute the JRE. The JRE consists of the Java Virtual Machine, the Java Core Classes, and supporting files. (Definition copyright 1996-1999 Sun Microsystems, Inc. All Rights Reserved. Used by permission.)

Java Server Page (JSP). Java Server Pages are Web pages that include dynamic tags which are executed on the server. JSPs are the presentation layer for Web-based applications built in Java.

Java Virtual Machine (JVM). A software implementation of a central processing unit (CPU) that runs compiled Java code (applets and applications).

JavaDoc. Sun's tool for generating HTML documentation on classes by extracting comments from the Java source code files.

JavaScript. A scripting language used within an HTML page. Superficially similar to Java but JavaScript scripts appear as text within the HTML page. Java applets, on the other hand, are programs written in the Java language and are called from within HTML pages or run as standalone applications.

JDBC (Java Database Connectivity). In the JDK, the specification that defines an API that enables programs to access databases that comply with this standard.

JIT. See *Just-In-Time compiler*.

JNDI. See *Java Naming and Directory Interface*.

JNI. See *Java Native Interface*.

JRE. See *Java Runtime Environment*.

Just-In-Time compiler (JIT). A platform-specific software compiler often contained within JVMs. JITs compile Java bytecodes on-the-fly into native machine instructions, thereby reducing the need for interpretation.

JVM. See *Java Virtual Machine*.

L

LAN. See *local area network*.

large object (LOB). See *LOB*.

licensed internal code (LIC). Microcode that IBM does not sell as part of a machine, but instead, licenses to the customer. LIC is implemented in a part of storage that is not addressable by user programs. Some IBM products use it to implement functions as an alternate to hard-wire circuitry.

linker. A computer program for creating load modules from one or more object modules or load modules by resolving cross references among the modules and, if necessary, adjusting addresses. In Java, the linker creates an executable from compiled classes.

load module. A program unit that is suitable for loading into main storage for execution. The output of a linkage editor.

LOB (large object). A sequence of bytes representing bit data, single-byte characters, double-byte characters, or a mixture of single- and double-byte characters. A LOB can be up to 2 GB -1 bytes in length. See also CLOB, DBCLOB.

local area network (LAN). A computer network located in a user's premises within a limited geographic area.

logical partition (LPAR). A set of functions that create a programming environment that is defined by the ESA/390 architecture. ESA/390 architecture uses this term when more than one LPAR is established on a processor. An LPAR is conceptually similar to a virtual machine environment except that the LPAR is a function of the processor. Also, LPAR does not depend on an operating system to create the virtual machine environment.

logical switch number (LSN). A two-digit number used by the I/O Configuration Program (IOCP) to identify a specific ESCON® Director.

logically partitioned (LPAR) mode. A central processor mode, available on the Configuration frame when using the PR/SM™ facility, that allows an operator to allocate processor hardware resources among logical partitions. Contrast with *basic mode*.

LPAR. See *logical partition*.

M

megabyte (MB). (1) For processor storage, real and virtual storage, and channel volume, 220 or 1 048 576 bytes. (2) For disk storage capacity and communications volumes, 1 000 000 bytes.

method. A fragment of Java code within a class that can be invoked and passed a set of parameters to perform a specific task.

middle tier. The hardware and software that resides between the client and the enterprise server resources and data. The software includes a Web server that receives requests from the client and invokes Java servlets to process these requests. The client communicates with the Web server via industry standard protocols such as HTTP and IIOP.

middleware. A layer of software that sits between a database client and a database server, making it easier for clients to connect to heterogeneous databases.

multithreading. Multiple TCBs executing one copy of code concurrently (sharing a processor) or in parallel (on separate central processors).

N

National Committee for Information Technology Standards. NCITS develops national standards, and its technical experts participate on behalf of the United States in the international standards activities of ISO/IEC JTC 1, information technology.

native class. Machine-dependent C code that can be invoked from Java. For multi-platform work, the native routines for each platform need to be implemented.

native SQL procedure. An SQL procedure that is processed by converting the procedural statements to a native representation that is stored in the database directory, as is done with other SQL statements. When a native SQL procedure is called, the native representation is loaded from the directory, and DB2 executes the procedure.

NCITS. See *National Committee for Information Technology Standards*.

NUL terminator. In C, the value that indicates the end of a string. For character strings, the NUL terminator is X'00'.

NUL-terminated host variable. A varying-length host variable in which the end of the data is indicated by the presence of a NUL terminator.

null. A special value that indicates the absence of information.

O

object. The principal building block of object-oriented programs. Objects are software programming modules. Each object is a programming unit consisting of related data and methods.

ODBC. See Open Database Connectivity.

ODBC driver. A dynamically-linked library (DLL) that implements ODBC function calls and interacts with a data source.

Open Database Connectivity (ODBC). A Microsoft® database application programming interface (API) for C that allows access to database management systems by using callable SQL. ODBC does not require the use of an SQL preprocessor. In addition, ODBC provides an architecture that lets users add modules called database drivers that link the application to their choice of database management systems at runtime. This means that applications no longer need to be directly linked to the modules of all the database management systems that are supported.

open system. A system whose characteristics comply with standards made available throughout the industry and that therefore can be connected to other systems complying with the same standards.

original equipment manufacturers information (OEMI). A reference to an IBM guideline for a computer peripheral interface. More specifically, refer to IBM S/360 and S/370 Channel to Control Unit Original Equipment Manufacturer's Information. The interface uses ESA/390 logical protocols over an I/O interface that configures attached units in a multi-drop bus environment.

P

package. A program element that contains classes and interfaces.

partition-by-growth table space. A universal table space whose size can grow to accommodate data growth. DB2 for z/OS manages partition-by-growth table spaces by automatically adding new data sets when the database needs more space to satisfy an insert operation.

partitioned by range table space. A type of universal table space that is based on user defined partitioning ranges.

persistence. In object models, a condition that allows instances of classes to be stored externally, for example in a relational database.

Persistence Builder. In VisualAge for Java, a persistence framework for object models, which enables the mapping of objects to information stored in relational databases and also provides linkages to legacy data on other systems.

plan. See *application plan*.

plan name. The name of an application plan.

precompilation. A processing of application programs containing SQL statements that takes place before compilation. SQL statements are replaced with statements that are recognized by the host language compiler. Output from this precompilation includes source code that can be submitted to the compiler and the database request module (DBRM) that is input to the bind process.

prepare. The first phase of a two-phase commit process in which all participants are requested to prepare for commit.

prepared SQL statement. A named object that is the executable form of an SQL statement that has been processed by the PREPARE statement.

primary key. A unique, non-null key that is part of the definition of a table. A table cannot be defined as a parent unless it has a unique key or primary key.

process. A program executing in its own address space, containing one or more threads.

program temporary fix (PTF). A temporary solution or bypass of a problem diagnosed by IBM in a current unaltered release of a program.

property. An initial setting or characteristic of a bean, for example, a name, font, text, or positional characteristic.

PTF. See *program temporary fix*.

R

RDBMS. Relational database management system.

reentrant. Executable code that can reside in storage as one shared copy for all threads. Reentrant code is not self-modifying and provides separate storage areas for each thread. Re-entrancy is a compiler and operating system concept, and re-entrancy alone is not enough to guarantee logically consistent results when multithreading.

reference. An object's address. In Java, objects are passed by reference rather than by value or by pointers.

relational database management system (RDBMS). A relational database manager that operates consistently across supported IBM systems.

remote. Refers to any object maintained by a remote DB2 subsystem; that is, by a DB2 subsystem other than the local one. A remote view, for instance, is a view maintained by a remote DB2 subsystem. Contrast with local.

Remote Method Invocation (RMI). RMI is a specific instance of the more general term RPC. RMI allows objects to be distributed over the network; that is, a Java program running on one computer can call the methods of an object running on another computer. RMI and java.net are the only 100% pure Java APIs for controlling Java objects in remote systems.

Remote Object Instance Manager. In Remote Method Invocation, a program that creates and manages instances of server beans through their associated server-side server proxies.

Remote Procedure Calls (RPC). RPC is a generic term referring to any of a series of protocols used to execute procedure calls or method calls across a network. RPC allows a program running on one computer to call the services of a program running on another computer.

requester. Also application requester (AR). The source of a request to a remote RDBMS, the system that requests the data.

RMI (Remote Method Invocation). See *Remote Method Invocation*.

rollback. The process of restoring data changed by SQL statements to the state at its last commit point. All locks are freed. Contrast with commit.

RPC. See *Remote Procedure Calls*.

runtime system. The software environment where compiled programs run. Each Java runtime system includes an implementation of the Java Virtual Machine.

S

sandbox. A restricted environment, provided by the Web browser, in which Java applets run. The sandbox offers them services and prevents them from doing anything naughty, such as doing file I/O or talking to strangers (servers other than the one from which the applet was loaded). The analogy of applets to children led to calling the environment in which they run the "sandbox."

scalar function. An SQL operation that produces a single value from another value and is expressed as a function name followed by a list of arguments enclosed in parentheses. See also column function.

Secure Socket Layer (SSL). SSL is a security protocol that allows communications between a browser and a server to be encrypted and secure. SSL prevents eavesdropping, tampering, or message forgery on your Internet or intranet network.

security. Features in Java that prevent applets downloaded off the Web from deliberately or inadvertently doing damage. One such feature is the digital signature, which ensures that an applet came unmodified from a reputable source.

serialization. Turning an object into a stream, and back again.

server. The computer that hosts the Web page that contains an applet. The .class files that make up the applet, and the HTML files that reference the applet reside on the server. When someone on the Internet connects to a Web page that contains an applet, the server delivers the .class files over the Internet to the client that made the request. The server is also known as the originating host.

server bean. The bean that is distributed using RMI services and is deployed on a server.

servlet. See *Java servlet*.

SGML. See *Standardized Generalized Markup Language*.

Shell. The user interface of UNIX system software. In z/OS, an xpg4.2-compliant shell is used. Very often OMVS is used as an interface for z/OS shells.

single precision. A floating-point number that contains 32 bits. See also double precision.

Small Computer System Interface (SCSI). (1) An ANSI standard for a logical interface to computer peripherals and for a computer peripheral interface. The interface uses a SCSI logical protocol over an I/O interface that configures attached targets and initiators in a multi-drop bus topology. (2) A standard hardware interface that enables a variety of peripheral devices to communicate with one another.

SmartGuide. In IBM software products, an active form of help that guides you through common tasks.

source type. An existing type that is used to internally represent a distinct type.

sourced function. A function that is implemented by another built-in or user-defined function already known to the database manager. This function can be a scalar function or a column (aggregating) function; it returns a single value from a set of values (for example, MAX or AVG). Contrast with external function and built-in function.

SQL. Structured Query Language. A language used by database engines and servers for data acquisition and definition.

SSL. See *secure socket layer*.

Standardized Generalized Markup Language. An ISO/ANSI/ECMA standard that specifies a way to annotate text documents with information about types of sections of a document.

static bind. A process by which SQL statements are bound after they have been precompiled. All static SQL statements are prepared for execution at the same time. Contrast with dynamic bind.

static SQL. SQL statements, embedded within a program, that are prepared during the program preparation process (before the program is executed). After being prepared, the SQL statement does not change (although values of host variables specified by the statement might change).

stored procedure. A user-written application program, that can be invoked through the use of the SQL CALL statement.

Structured Query Language (SQL). A standardized language for defining and manipulating data in a relational database.

Sysout. The regular output for a program on z/OS is SYSOUT. It is the functional equivalent of stdout on UNIX. In batch, there can be multiple SYSOUTs.

System. A single instance of the z/OS or OS/390 operating system in a sysplex.

System Management End User Interface (SMEUI). A Windows-based tool that makes it possible to perform administrative tasks for WebSphere Application Server from a Windows workstation. The SMEUI tool is used to deploy a new application to WebSphere on z/OS.

T

task control block (TCB). A control block used to communicate information about tasks within an address space that are connected to DB2. An address space can support many task connections (as many as one per task), but only one address space connection. See *address space connection*.

TCB. Task Control Block; manages dispatchable tasks. Each UNIX thread is assigned to a TCB.

Telnet. Telnet provides a virtual terminal facility that allows users of one computer to act as if they were using a terminal connected to another computer. The Telnet client program communicates with the Telnet daemon on the target system to provide the connection and session.

temporary table. A table created by the SQL CREATE GLOBAL TEMPORARY TABLE statement that is used to hold temporary data. Contrast with result table.

textual XML format. A system of storing XML data in text, as opposed to binary, that allows for direct human reading.

thin client. Thin client usually refers to a system that runs on a resource-constrained machine or that runs a small operating system. Thin clients don't require local system administration, and they execute Java applications delivered over the network.

third tier. The third tier, or back end, is the hardware and software that provides database and transactional services. These back-end services are accessed through connectors between the middle-tier Web server and the third-tier server. Though this conceptual model depicts the second and third tier as two separate machines, the NCF model supports a logical three-tier implementation in which the software on the middle and third tier is on the same box.

thread. A separate flow of control within a program.

timestamp. A seven-part value that consists of a date and time expressed in years, months, days, hours, minutes, seconds, and microseconds.

trace. A facility that provides the ability to monitor and collect monitoring, auditing, performance, accounting, statistics, and serviceability data.

Trading communities. Trading communities bring together buyers and sellers in a central online location to trade, using various online mechanisms including auctions and exchanges, in addition to industry content and application services. Trading communities are owned and operated by both large industry players in closed trading networks and by neutral parties in more fragmented open communities.

transaction. (1) In a CICS program, an event that queries or modifies a database that resides on a CICS server. (2) In the Persistence Builder, a representation of a path of code execution. (3) The code activity necessary to manipulate a persistent object. For example, a bank application might have a transaction that updates a company account.

U

UDF. See *user-defined function*.

UDT. See *user-defined data type*.

Unicode. A 16-bit international character set defined by ISO 10646. See also *ASCII*.

Uniform Resource Locator (URL). The unique address that tells a browser how to find a specific Web page or file.

universal table space. A table space that is both segmented and partitioned.

URI/URL. A Uniform Resource Identifier (URI) and Uniform Resource Locator (URL) uniquely define a location on the Web. URLs are familiar to anyone who browses the Web (for example <http://www.ibm.com>), and the term URI is a more general term which also incorporates other schemes for identifying resources.

URL. See *Uniform Resource Locator*.

user-defined data type (UDT). See *distinct type*.

user-defined function (UDF). A function defined to DB2 using the CREATE FUNCTION statement that can be referenced thereafter in SQL statements. A user-defined function can be either an external function or a sourced function. Contrast with built-in function.

V

valid. An XML document is valid if its content conforms to the rules in its DTD.

variable. (1) An identifier that represents a data item whose value can be changed while the program is running. The values of a variable are restricted to a certain data type. (2) A data element that specifies a value that can be changed. A COBOL elementary data item is an example of a variable. Contrast with constant.

vi. A popular UNIX editor. It can only be used from an ASCII Telnet connection.

virtual machine. A software or hardware implementation of a central processing unit (CPU) that manages the resources of a machine and can run compiled code. See *Java Virtual Machine*.

visual bean. In the Visual Composition Editor, a bean that is visible to the end user in the graphical user interface.

W

WAP. Wireless Application Protocol. Offers Internet browsing from wireless handsets.

Web. See *World Wide Web*.

Web Application. A WebSphere Web application is a collection of static pages, JSPs, and Servlets that share a common URL prefix, and together make a complete application.

Web browser. The Web uses a client/server processing model. The Web browser is the client component. Examples of Web browsers include Mosaic, Netscape Navigator, and Microsoft Internet Explorer. The Web browser is responsible for formatting and displaying information, interacting with the user, and invoking external functions, such as Telnet, or external viewers for data types that it does not directly support. Web browsers are fast becoming the universal client for the GUI workstation environment, in much the same way that the ability to emulate popular terminals such as the DEC VT100 or IBM 3270 allows connectivity and access to character-based applications on a wide variety of computers. Web browsers are available for all popular GUI workstation platforms and are inexpensive (often included with operating systems or related products for no additional charge.)

Web server. Web servers are responsible for servicing requests for information from Web browsers. The information can be a file retrieved from the server's local disk or generated by a program called by the server to perform a specific application function. Web servers are sometimes referred to as httpd servers or daemons. A number of Web servers are available for most platforms including most UNIX variants, OS/2 Warp, OS/390, and Windows NT®.

well-formed. An XML document is well-formed if there is one root element, and all its child elements are properly nested within each other. Start tags must have end tags, and each empty tag must be designated as such with a trailing slash. Also, all attributes must be quoted, and all entities must be declared.

white-space. In XML, characters that are not visible, but used in formatting documents or programs. These characters include the SPACE, TAB, NEWLINE, and CARRIAGE-RETURN characters.

World Wide Web. A network of servers that contain programs and files. Many of the files contain hypertext links to other documents available through the network.

WWW. See *World Wide Web*.

X

XML. The Extensible Markup Language (XML) is an important new standard emerging for structured documents on the Web. XML extends HTML beyond a limited tag set and adapts SGML, making it easy for developers to write programs that process this markup and providing for a rich, more complex encoding of information.

XML attribute. A name-value pair within a tagged XML element that modifies certain features of the element.

XML column. A column of a table that stores XML values and is defined using the data type XML. The XML values that are stored in XML columns are internal representations of well-formed XML documents.

XML data type. A data type for XML values.

XML element. A logical structure in an XML document that is delimited by a start and an end tag. Anything between the start tag and the end tag is the content of the element.

XML index. An index on an XML column that provides efficient access to nodes within an XML document by providing index keys that are based on XML patterns.

XML lock. A column-level lock for XML data. The operation of XML locks is similar to the operation of LOB locks.

XML node. The smallest unit of valid, complete structure in a document. For example, a node can represent an element, an attribute, or a text string.

XML node ID index. An implicitly created index, on an XML table that provides efficient access to XML documents and navigation among multiple XML data rows in the same document.

XML pattern. A slash-separated list of element names, an optional attribute name (at the end), or kind tests, that describe a path within an XML document in an XML column. The pattern is a restrictive form of path expressions, and it selects nodes that match the specifications. XML patterns are specified to create indexes on XML columns in a database.

XML publishing function. A function that returns an XML value from SQL values. An XML publishing function is also known as an XML constructor

XML schema. In XML, a mechanism for describing and constraining the content of XML files by indicating which elements are allowed and in which combinations. XML schemas are an alternative to document type definitions (DTDs) and can be used to extend functionality in the areas of data typing, inheritance, and presentation.

XML schema repository (XSR). A repository that allows the DB2 database system to store XML schemas. When registered with the XSR, these objects have a unique identifier and can be used to validate XML instance documents.

XML serialization function. A function that returns a serialized XML string from an XML value.

XML table. An auxiliary table that is implicitly created when an XML column is added to a base table. This table stores the XML data, and the column in the base table points to it.

XML table space. A table space that is implicitly created when an XML column is added to a base table. The table space stores the XML table. If the base table is partitioned, one partitioned table space exists for each XML column of data.

XSL Stylesheet. The eXtensible Stylesheet Language defines stylesheets for XML Documents. It is composed of two parts: the formatting objects, and XSLT. XSL is defined by the WorldWide Web Consortium.

XSLT. eXtensible Stylesheet Language Transformations. This defines the part of the XSL specification which allows the stylesheet to reformat and reorganize the XML data. It is most often used to transform XML into XSL.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 294. Note that some of the documents referenced here might be available in softcopy only.

- ▶ *DB2 10 for z/OS Technical Overview*, SG24-7892
- ▶ *DB2 9 for z/OS Stored Procedures: Through the CALL and Beyond*, SG24-7604
- ▶ *XML on z/OS and OS/390: Introduction to a Service-Oriented Architecture*, SG24-6826
- ▶ *XML Processing on z/OS*, SG24-7810

Other publications

These publications are also relevant as further information sources:

- ▶ *DB2 10 for z/OS Installation and Migration Guide*, GC219-2974
- ▶ *DB2 10 for z/OS pureXML Guide*, SC19-2981
- ▶ *DB2 10 for z/OS Application Programming Guide and Reference for Java*, SC19-2970
- ▶ *DB2 10 for z/OS Application Programming and SQL Guide*, SC19-2969
- ▶ *DB2 10 for z/OS SQL Reference*, SC19-2983
- ▶ *DB2 pureXML Cookbook*, Matthias Nicola and PAV Kumar-Chatterjee, IBM Press, ISBN-13: 978-0-13-815047-1
- ▶ *WebSphere MQ Application Programming Guide Version 6.0*, SC34-6595
- ▶ *WebSphere MQ Using Java Version 6.0*, SC34-6591

Online resources

These Web sites are also relevant as further information sources:

- ▶ Tools and XML functionality for DB2 pureXML users
<https://www.ibm.com/developerworks/data/library/techarticle/dm-1012xmltools/>
- ▶ Extensible Dynamic Binary XML, Client/Server Binary XML Format(XDBX) Version 1.0
<http://www.ibm.com/support/docview.wss?uid=swg27019354&aid=1>
- ▶ XSL Transformations (XSLT) Version 1.0
<http://www.w3.org/TR/xslt>
- ▶ pureXML Devotees
<http://www.ibm.com/developerworks/wikis/display/db2xml/devotee#devotee-rational>

- ▶ ISO 3166 code lists
http://www.iso.org/iso/country_codes/iso_3166_code_lists.htm
- ▶ OASIS
http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=ubl
- ▶ ISO 20022 Universal financial industry message scheme
<http://www.iso20022.org/>
- ▶ Catalogue of ISO 20022 messages
http://www.iso20022.org/catalogue_of_unifi_messages.page

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Numerics

2-byte length 195
variable character 220

A

access plan 28
ALTER statement 75
ALTER TABLE 54, 56, 74–76, 175, 232
statement 56, 74, 77
Amt Ccy 68, 265
APIs 129–130, 132
application xix, 4–5, 19–20, 22–23, 31, 45–46, 52, 59, 67, 70, 88, 93, 129–130, 155–156, 213, 235, 243–244, 246, 274
application scenario 118, 246
applications with SQL (AS) 184
array 131
attribute 6, 9, 27, 34, 36, 78, 80, 100, 132, 157, 174, 176, 191, 234, 247, 250
auxiliary index 184
auxiliary table 62, 257
availability 3, 48, 58

B

bank statement 102–103, 171–172, 234
base table 52–53, 58, 88, 103, 164, 182–184, 231–232, 258, 261
document ID column 61
non-XML columns 197
page size 233
partition-by-growth table space 66
row changes partition 56
table space 54
XML indicator column 54
base table space 53, 184
BKRTORCS 63, 240
DSN00242.BKRT ORCS 190
set 238
binary format 135, 158, 199
BK_TO_CSTMR_STMT c 116, 253
BK_TO_CSTMR_STMT Position 196
BLOB 98, 132, 156, 198
BLOBs 156, 159
buffer 63, 230–231, 258–259
Byte Order Mark (BOM) 157

C

C 26, 28, 85, 216–217, 226–227, 230, 253–255
c.xml address 26
CCSID 157, 198
cdc message 118–120
potential use 121

change data 49, 118–119
XML 118
CHECK DATA 182–183, 237
CHECK Data 182, 237
default behavior 182
CHECK INDEX 186, 239
CHECK Index 187, 239
class 4, 22, 133, 186, 275
CLASSPATH 130
client application 135
CLOB 31, 91, 94–96, 132, 156–157, 190, 198, 208, 211, 226
CLOBs 156, 159
COBOL xix, 19, 22, 45–46, 48–49, 87, 155–156, 198, 245, 256–257, 273, 276
COBOL program 50, 173, 175, 276
code page 156–157, 256–257
code page conversion 177
Column
XML 23, 65, 103, 158, 184
Command Line Processor 37, 89, 91, 138, 163, 235
Command Line Processor (CLP) 37, 89
Comments xxi, 7, 44, 119
COMMIT 41, 59, 97, 248
complete xmlschema
SYSXSR.MYXMLSCHEMA 37
components 137, 158, 161, 173
compression 61, 234, 258
condition 92–93
connection 100, 136, 229
constraint 34, 64
constructor 107–108
Content 3, 46, 74, 178, 207, 250
Context 7, 27, 57, 107, 170
COPY utility
control statement 188
COPYTOCOPY 191
corresponding namespace name
schema location hint 80
Create 4, 24–25, 47, 51, 74, 88, 107, 130, 133, 160, 192, 195–196, 230, 236, 244, 251, 274
CREATE PROCEDURE 89, 91, 96
CreDtTm element 170
MsgRcpt element 175
MsgRcpt element right 175
cross-loader 192
CURRENT TIMESTAMP 92

D

data access 22–24, 260
Data compression 258
data element 23, 30, 88, 101, 245
large number 31
data format 136

data model 23, 31–32, 107, 126, 246
 data set 41, 58, 63, 195–197, 199, 230, 246
 data source 120
 data structure 12, 22, 126, 246
 Data Studio 44, 93
 data type xix, 10–11, 19, 22–23, 35, 52, 54, 56, 74–75, 88, 90, 111, 113, 130, 156, 159, 196, 220, 236–237, 240, 244, 249–250
 BIGINT 61
 casing 35
 casting 29
 timestamp 91
 XML 23, 54, 74–76, 88, 90, 130, 156, 159, 175, 220, 249–250
 database access thread (DBAT) 42
 DATABASE DSN00242 67, 193
 database objects 246
 databases 2, 22–23, 129, 275
 DB0B DSNTDDIS 67, 207, 239
 DB2 10 21, 74, 88, 167, 183, 230, 257
 MQ functions 94
 online compression 258
 zAAP eligible 42
 DB2 9 xix–xx, 23, 42, 84, 93, 167, 226, 231, 256
 DB2 data 23, 38, 132, 158
 DB2 database xx, 22–23, 38, 52, 260
 DB2 engine 42, 49, 256
 DB2 for z/OS xix, 19, 42, 119, 129–130, 163, 192, 237, 276
 DB2 pureXML
 Using COBOL 50
 XML cdc messages 121
 DB2 subsystem 88, 94, 97, 231
 DB2 table 5, 24, 29, 32, 88, 90–91, 248, 275
 XML document 32
 db2 table 91, 248
 DB2 V8 23
 DB2 z/OS 257
 DB2Connection 235
 DBA 28, 229
 DBCLOB 156, 198
 DBD Length 67, 207, 239
 DD DISP 165
 DD DSN 186, 259
 DDL 48, 155, 233
 DE LETE 187
 default element namespace 15, 28, 90, 106, 169, 208, 236, 251
 default namespace 15, 107, 170
 default value 197
 DELETE 5, 33, 62, 168, 186–187, 240, 258–259
 delete 5, 32–33, 115, 167, 186, 215, 240–241, 259
 DEV Type 216
 DOCID 54, 183, 246
 DOCID column 54, 61, 183
 corresponding value 183
 unique index 54
 DOCID index 54, 205, 214, 239, 253
 DOCID key 239
 DOCTYPE 5, 14, 16

DOCTYPE declaration 16
 Document Object Models 7
 Document Type Definition 6–7
 Document validity 5
 Document xmlns
 xsi 197, 199, 269
 Documentation 11, 47, 93
 DOM 7, 129, 133–135
 See also also also also Document Object Model
 See also also also Document Object Model
 See also also Document Object Model
 DRDA 42, 192
 DSN_XMLVALIDATE function 34, 93
 DSN_XMLVALIDATE invocation 42, 82
 DSN1COPY 228
 DSNE610I Number 63, 166, 190, 252
 DSNE616I STATEMENT Execution 63, 166, 204, 252
 DSNT408I SQLCODE 166, 169
 DSNT415I SQLERRP 166, 169
 DSNT416I SQLERRD 166, 169
 DSNT418I SQLSTATE 166, 169
 DSNTIAUL 225
 DSNUPROC.SYSIN DD 259
 DTD 5–7
 DtTm element 70, 108, 173
 dynamic 75, 130, 192, 220

E

EBCDIC CCSID 157, 200
 element 6, 23, 26–27, 48, 78, 90, 134, 162, 190, 208, 235–236, 247, 250–251, 264, 276–277
 element name 7, 9, 36, 78, 112, 139, 170
 email emailUse 31
 encoding scheme 54, 56, 60
 Enterprise Service Bus 45–46, 88, 93
 ENVIRONMENT 15, 38, 40, 94, 137, 176, 216–217, 249, 274
 environment 19, 38–39, 93–94, 97, 129, 137, 162, 230, 249, 256, 275
 eq 15
 EXEC SQL 59, 157, 165, 192–193
 OPEN CURSOR C1 59
 statement 171
 EXPLAIN 7, 21, 90, 173, 251, 254–255
 expression 14, 27–28, 82, 103, 107–108, 169, 171, 173, 236, 250, 252
 eXtensible Markup Language 4
 Extensible Stylesheet Language 11
 extensions 24, 107

F

FETCH 59, 172
 fetch 123
 file reference variable 160, 197–198, 222, 276
 variable declarations 160
 fn
 abs 107
 empty 108
 fragment 17, 33

function 21, 25–26, 28, 50, 74, 81–82, 90–91, 130, 139, 163, 165–167, 192, 214, 248, 254
function DSN_XMLVALIDATE 42

G

GENERATED ALWAYS 54, 61
given table space
 XML data 224
GrpHdr element 166
 Schema definition 170
GUI 43

H

handle 2, 51, 56, 99, 120, 129, 158–159, 175, 181, 225, 237
Hierarchical xix, 7, 32, 247
HIGH DSNUM 216
host variable 157–159
host variables 156
HTML xxi, 4, 31, 137, 150, 276

I

IC Type 216
ID attributes 15
ID INT 76
II14426 38, 231, 261
image copy 188
import 48, 134
INCLUDE XML TABLESPACES 183–184, 239
INDDN SYSREC00 226
index access 30, 109, 111, 173, 253–255
INDEX XMLR4 197, 199
indexes 22–23, 28, 47, 51, 58, 104, 164, 172, 186, 230, 244, 246
Information Integrator 192
input data 176, 195–196, 198, 237
input parameter 88–90, 139, 165
INSERT 24–25, 31, 51, 74, 77, 80, 90, 92, 132, 156, 165, 168, 196, 234–235, 253, 275
installation 21, 38, 94, 230
Installation job
 DSNTIJMV 39
 DSNTIJRT 38, 230
 DSNTIJRW install 38–39
instance document 11, 80
 root element 80
 root element node 83
Interactive Financial Exchange (IFX) 2
IS xix, 2, 21, 45, 52, 63, 74, 87, 130, 156, 182, 187, 229, 244, 252, 263, 273
ISO/IEC 24
ISO20022 standard 48, 118, 155, 246
 BankToCustomerStatement message 161
 same subset 161

J

jar 130
Java xix, 19, 22, 38–39, 45, 48, 50, 91, 129, 230, 235,

245, 257, 273–275
java program 46, 93
JCL 97, 164–165, 180, 186, 188, 190–191, 193, 258, 277
JCLLIB Order 186
JDBC 22, 38–39, 41, 50, 129–130, 230, 235, 257, 274–275
 driver 130
JOBPARM SYSAFF 186
John Doe 26, 121
joins 116

K

KB 195, 231, 259, 274
Key 15, 35, 54, 57, 61, 106, 165, 234, 251, 258
keyword 8, 77–78, 177, 182–184, 238–239, 258

L

LANGUAGE SQL 90, 92, 96
let 17–18, 27, 74, 89, 99, 108, 159, 162, 204, 233
LISTDEF 193, 238
LISTDEF List 193
LISTDEF LISTALL 193
LISTDEF LISTXML 193
LISTDEF utility 193
LOAD 70, 77, 107, 135, 166, 192, 195, 235
LOAD PHASE Statistic 197
LOAD utility 196, 198, 237
 crossloader capability 237
 XML data 197
LOB 52, 156, 181–182, 184, 237
LOB table 184, 238
LOBs 156, 160, 184, 227, 237
location 13, 26–27, 77, 116, 118, 138, 166, 192, 236, 250–251
locking 32, 62
LOW DSNUM 216
LRSN 60, 204–205, 208

M

markup language 4
MERGE 202–203, 235
MERGECOPY 202
message queue
 XML message 98
message queue (MQ) 93
metadata 4
method 9, 116, 131, 173, 197, 225–226, 235–236, 248, 258
monitor 249, 259
Move XML-Text 178
MQ Listener
 configuration data 98
 process 97
MQ listener 93
 configuration 97
MSG_CRE_DT_TM TIMESTAMP 88, 192
MSG_ID VARCHAR 224
MsgRcpt element 167, 277

multiple occurrences 174
XML document 169

N

name space 8
namespace 8, 27–28, 75, 77, 90, 162, 168, 190, 208, 235–236, 251
namespace declaration 15, 28, 169, 253
 MsgRcpt element 170
namespace name 80
 schema location hint 80
namespace prefix 8–9, 107
Namespaces 8
namespaces 5, 7–9, 27, 35, 85, 103, 170, 246
native data type 24
native SQL 49, 88, 256
Node 14, 32, 57, 80, 103, 169
node 14, 30, 32–33, 54, 56–57, 80, 101, 103, 134–135, 168–169, 173, 190, 258, 276
node id index 61
NODEID index
 corresponding entry 183
 index entry 183
 XML table space 188
NORMAL Completion 67, 207, 239
Ntry element 101
NUMRECS 1 200

O

ODBC 39
optimizer 28, 35, 249, 251
options 11, 120, 156, 159, 163, 183–184, 189, 211, 240, 257
order by 59, 65, 103–104, 123, 257
OUTPUT Start 187–188
overhead 258

P

package 41
parameter 41, 82, 88–89, 139, 165, 171, 225–227, 231, 259, 275
parent/child relationships 6, 18
parser 5, 7, 139–140, 178
partitioning 56, 61, 234, 258
PATH xix, 11, 28, 57, 92, 165, 169, 249, 254
Pattern 13, 34, 70, 109, 173, 250
PCDATA 6–7
performance xix, 19–20, 23, 34, 58, 131, 137, 224, 231, 236, 243–244
persisting 244
PHASE Statistic 197
PIT LRSN 216
PK90032 42
PK90040 42
PM21834 240
PM22081 235
PM24947 240
PM26592 240

PM28385 125
PM29986 201
po
 purchaseOrder xmlns
 po 80
point-in-time recovery 206, 238
precompiler 156
predicate 26–27, 103, 108, 171–172, 253–254
prefix 8–9, 107, 170, 226
primary schema 38
privilege 236–237
procedures 24, 37, 46, 48–49, 87, 138, 159, 230, 245–246, 248, 273
PROCESSING SYSIN 187
programming interface 126
programming language 130–131
Publishing xix, 24–25, 89, 124, 133
purchase order 3, 75
PURCHASE_ORDERS Value 81
pureQuery 44
pureXML xix–xx, 19–22, 46–48, 87–88, 106, 137, 155–156, 161, 229–230, 235, 243, 245–246
pureXML storage 121, 246, 248

Q

qualified name 84
query 15, 26–28, 45–47, 63, 84, 93, 99, 137, 175, 179, 197, 207, 226, 234, 238, 246–247, 251
Querying XML documents 171
Queue Manager
 MQBA 97
QUIESCE 203

R

RACF 213–214
RBA 60, 204–205, 208
REBUILD INDEX 188, 205, 240
received XML document
 new XML document 91
 relational and XML objects 112
RECOVER 188, 206
Redbooks Web site 48, 273, 294
 Contact us xxi
REGION 186, 188, 190–191
registered XML schema
 XML documents 36
relational column 29, 49, 96, 115, 245
 data elements 246
relational data xix, 4, 23–24, 45, 90, 116, 176, 178, 251
relational index 34, 70, 236, 249
reordered row format 61
REORG 58, 188, 211, 234, 237, 258
REORG TABLESPACE 211, 237, 259
REORG UNLOAD ONLY 212
REORG utility 58, 197
REPORT 39, 137, 166, 182, 197, 215, 238
REPORT RECOVERY 216
 TABLESPACE DSN00242.BKRT ORCS 216
 TABLESPACE DSN00242.XBKR 0000 217

REPORT TABLESPACESET 215–216
 repository 22, 30, 36, 48, 78, 82, 96, 127, 138, 163, 175, 230, 276
 requirements 45–47, 74, 121, 273
 result set 103, 197
 ResultSet 131
 retrieved XML document
 XML structure 24
 return 15, 29–30, 89, 92, 99, 132, 166, 200, 225–226, 253, 255
 RETURN Code (RC) 39, 197, 199
 RETURNS VARCHAR 41, 113
 root element 11, 75
 namespace name 80
 node 80
 row format 61
 RUNSTATS 106, 173, 219, 237, 252
 runtime 27

S

samples xix, 19, 48, 137, 273
 SAX 129, 132, 137
 Parser 132
 Schema 5–6, 34, 38, 44, 47, 50, 74, 79, 89, 138, 162–163, 179, 184, 230, 236, 247–248, 264, 275–277
 XML 6, 35, 38, 44, 74, 162, 230, 248
 schema 6, 21–22, 27, 48, 74, 89, 137–138, 155, 161–162, 182–183, 230, 246–247, 264, 274–275
 schema document 38, 78, 80, 247
 schema location
 hint 80
 hint http 81
 URI 80
 schema validation 22, 34–35, 48, 74, 78, 91, 95–96, 107, 139, 163–164, 179–180, 185, 230, 248
 Schemas 7, 9, 22, 47, 74, 79, 89, 130, 162, 230, 244, 264
 SCOPE PENDING 182
 SCOPE XMLSCHEMAONLY 183
 scripts 155
 SDK 137, 275
 SELECT statement 110, 226–227, 260
 SELECT SUBSTR 63
 SELECT XMLSERIALIZE 208
 Semi-structured data 247
 separate file 7, 195, 220
 Server 38–39, 41, 119, 130, 192, 220, 230, 257, 273–275
 SET 6, 22, 33, 47, 63, 74, 76–77, 90, 130, 165, 168, 190, 192, 195, 230, 241, 249, 254
 setup 40–41, 93–94, 137, 230, 273–274
 SGML 4
 SHR LVL 216
 Shredding 178–179
 shredding 49, 139, 161, 178
 SHRLEVEL 185, 235
 side 137, 175–176
 SOA 2, 44
 source file 18
 spanned record format 195, 201
 SQL xx, 12, 19, 22–24, 45–47, 49, 51, 59, 67, 77–78, 84,

87–88, 130, 132, 156–157, 192, 231, 234–236, 243, 245, 275
 SQL code 108
 SQL error 158, 252
 SQL PL 62
 SQL programmer 112, 252
 SQL statement 24, 30, 37, 40, 84, 95, 100, 115, 168, 228, 256
 SQL Type 160
 SQL/XML 23–24, 47, 49, 87, 93, 99, 133, 159, 245–246
 SQL/XML extension 27
 SQL/XML language 260
 SQL/XML query 99, 253
 numbered points 103
 SQLCA 168, 171
 SQLCODE 41, 63–64, 90, 92, 166, 169, 204, 252
 SQLCODE Integer 96
 SQLJ 38–39, 41, 50, 129–130, 230, 257, 274–275
 sqlj 130
 SQLSTATE 92, 166, 169
 Standard Generalized Markup Language 4
 START LRSN 216
 startup procedure 38, 230
 statement 22, 27, 45–46, 48, 51–52, 54, 75, 89–90, 137, 161, 164, 184–186, 232, 252, 263, 276
 statistics 197, 211, 219, 237, 252
 STEPLIB 41, 165, 230
 storage model 244–245
 stored procedure 38, 45–46, 48, 88–89, 129, 139, 230, 256, 274
 string value 18, 27, 31–32, 133, 249
 stylesheet 11–12
 SYSIBM.SYSD UMMY1 41
 sysibm.sysdummy1 41, 90, 260
 SYSIBM.SYSI NDEXES 65, 204, 251
 SYSIN DD 165, 225

T

TABLE BK_TO_CSTMTR_STMT
 table space names 63
 table BK_TO_CSTMTR_STMT 52, 88, 164, 274
 INDDN SYSREC00 226
 table BK_TO_CUSTMR_STSMT 276
 TABLE PURCHASEORDERS 76
 table space 38, 52, 75, 186, 231, 246
 ACHKP status 241
 different table space 52
 logging attribute 191
 partial recovery 218
 SHRLEVEL CHANGE 212
 table T1 58
 tables 24, 31, 45, 47, 51, 60, 94, 115, 130, 162, 164, 182, 192–193, 230, 244, 247, 274
 TABLESPACE DSN00242.BKRT ORCS 187
 PARTITION 1 200
 TOLOGPOINT X'000011112222 208
 TOLOGPOINT X'000069667C5E 208
 TABLESPACE DSN00242.XBKR 0000 187
 COPYDDN 188
 PARTITION 1 200

target namespace 11, 77
 XML schema 77
 XML schemas 81
 task control block (TCB) 42
 text node 135
 three-layer structure 159
 Data conversion 159
 TIMESTAMP Path 92, 260
 tree structure 7, 12

U

UDF 42, 114–115
 UDFs 48, 112–113
 UK62510 240
 Unicode 2, 54, 56, 60, 156–157, 198
 Uniform Resource Identifier 8
 Uniform Resource Identifier (URI) 80
 universal table space 32, 52, 54, 58, 167, 211, 257
 UNLOAD 135, 187, 223, 237
 UPDATE 32–34, 50, 54, 59, 74, 77, 84, 90, 115, 124, 132, 165, 168–169, 190, 219, 258
 URI 8, 77–78, 236
 URL 13, 44
 USAGE 32, 46, 58, 106, 157–158, 247
 UTF-8 10, 31, 36, 54, 56, 60, 67, 120, 122, 131–132, 156–157, 197, 199, 256–257, 264, 269
 UTIL EXEC DSNUPROC 188
 UTILITY Execution 197, 199
 utility run 186, 193

V

V_CREDITM TIMESTAMP 89
 V_MSG_ID VARCHAR 92
 Validation 21, 48, 73–74, 77, 88, 139, 162, 179, 184, 230, 239, 248
 validity 5–6
 VALUE 2, 26, 46, 52, 74, 106, 115, 131, 164, 168, 183, 213, 239, 248
 VALUES 6, 24–25, 67, 75, 80–81, 90, 94, 129, 160, 165, 195, 213, 231, 248, 264
 VARCHAR 25, 27, 29, 52, 54, 56, 65, 82, 88–90, 139, 159, 164–165, 192, 196–198, 202, 233, 235–236, 251, 254–255
 VARCHAR NULLIF 198
 variable 96, 115–116, 130, 156, 195, 197–198, 246, 250, 276
 VBS data 224
 versions 22, 32, 48, 54, 58, 74, 83, 88, 119, 130, 211, 236–237, 258
 views 18

W

W3C 4, 9–10, 134
 Web browser 5, 273
 Web services 2–3, 44, 247
 WebSphere 14, 45, 87, 93–94, 158, 248, 276
 WebSphere Message Broker 118
 WebSphere MQ 46, 49–50, 88, 93–94, 276

well-formed XML 9, 30, 74
 Well-formedness 6
 well-formedness 5–7
 whitespace 31–32, 74, 196
 wide range 44, 87, 236, 247
 WLM 38, 94, 230
 WLM environment 39

X

XHTML 15–16
 XHTML 1.0 Frameset 16
 XHTML 1.0 Strict 16
 XHTML 1.0 Transitional 16
 XHTML example 17
 XID XMLADDRESS 26
 XLink 17
 XML xix–xx, 1–3, 19, 21–22, 29, 45–46, 51–52, 73–74, 87–88, 108, 116–118, 129–130, 155–156, 181–182, 229–230, 244, 263–264, 274–276
 definition 4, 6, 64, 74–75, 106, 139, 174, 222, 252
 editor 44
 Repository 22, 37, 79
 Schema 5, 9, 36, 90, 248
 standards 2, 24, 48, 248, 263
 Web Services 44
 XML column 24–27, 29, 37, 51–52, 54, 61, 74, 115, 129–131, 158, 163–164, 166–167, 174, 183, 185, 192, 195, 231–232, 246, 274, 276
 base table spaces 211
 billing statements 75
 data 26, 52, 57, 74, 130–131, 174, 232, 237
 definition 75
 document 74
 length 195
 need 76
 purchase orders 75
 Reset XML type modifier 77
 table space 215
 type modifier 174
 value 195, 220
 XML document 67
 XML table 57
 XML type identifier 180
 XML type modifier 76
 XML columns 24, 43, 52, 54, 58, 60–61, 75, 106, 115, 129, 132, 167, 182, 188, 192, 231–232, 244, 246
 XML data xix, 2, 21–22, 32, 46–47, 49, 52, 56, 60–61, 74, 88, 91, 103, 107, 129–130, 156, 158, 181–182, 186–187, 220, 223, 231–232, 234, 245–246, 248, 275
 file reference variables 160
 integrity rules 248
 internally encoded variable 159
 NODEID index 57
 required transformation 11
 Storage structure 56
 Using non-XML variables 159
 XML variables 159
 XML data model 36, 247
 XML data type 23, 74, 93, 95–96, 130, 132, 156, 159, 249

- XML declaration 131–132, 157
- XML declarations 132
- XML document 5, 22–23, 46, 50–51, 54, 56, 74, 80–81, 84, 88–89, 94, 112, 115, 129–130, 157, 165, 182–183, 189–190, 230, 233, 237, 245–246, 255, 275–276
 - Account Name 116
 - adhere 6
 - code page 158
 - credit transactions 102
 - data element 27
 - data elements 23
 - data values 258
 - declaration 11
 - detailed constraints 6
 - efficient insert 160
 - exact size 197
 - first few characters 197
 - Generation 176
 - individual elements 29
 - internal structure 23
 - internal structures 17
 - large number 104
 - location 237 169
 - multiple copies 257
 - multiple elements 93
 - multiple parts 70
 - multiple versions 54
 - new version 32
 - new XML node 115
 - old version 32
 - relational format 172
 - root element 11, 80
 - similar support 178
 - single element 101
 - source tree 12
 - specific fields 96
 - structural integrity 183
 - syntactical rules 6
 - textual nature 258
 - transactional entry 103
 - type 16, 246, 249
 - update 257
 - Validation 176
 - validation 84, 90–91, 93, 175, 249
 - viewer 100
 - XML data elements 107
 - XML elements 28
- XML element 29, 34, 167, 169, 177, 250
- XML Extender 23
- XML file 7, 12, 17, 156, 160, 164, 197, 226, 276–277
- XML format 50, 102, 117, 131, 135, 173, 196, 220, 237, 257
- XML Index
 - pattern 251
- XML index 28, 34–35, 70, 104–107, 172–173, 188, 196, 231, 235–236, 246, 249–250
 - CHECK INDEX 188
 - Creation 236
 - following keywords 219
 - XML pattern 253
- XML indicator column 54, 61
- XML message 46, 48, 50, 88, 95–96, 137, 167, 264, 274, 276
 - common source 118
 - existing standard 47
- XML model 22, 107, 248
- XML namespaces 35
- XML object 66, 183, 215, 231
 - Backup and recovery 237
 - following items 183
 - REPAIR utility 215
- XML PARSE
 - NEW-RCPT 178
 - RCPT 179
 - statement 178–179
- XML parser 7
- XML parsers 6
- XML pattern 35, 70, 250
 - wide range 251
- XML patterns 247, 250–251
- XML processing 91, 230, 249
 - error handling 91
- XML record 60, 197
 - logical creation 60
 - logical deletion 60
- XML Schema 5–6, 36, 74, 120, 179, 230, 243, 250
 - Definition 9
 - Language 11
- XML schema 6, 22, 27, 34–36, 48, 74, 89, 94–96, 106, 125, 138, 162–166, 183, 230, 235, 246–250, 275–276
 - naming standard 236
 - optional schema location 82
 - registration process 38
 - XML documents 179
- XML schema repository 37, 82, 230, 236
- XML schema validation 35, 37, 74, 82, 248
- XML Schemas 11, 22, 48, 74, 106, 138, 163, 236, 246
- XML source
 - document 12
- XML structure 22, 91, 246
 - inherent strengths 246
- XML table 52, 57–58, 75, 216, 219, 231, 233, 249
 - MIN_NODEID column 57
 - new version 59
 - NODEID index 239
 - space DSN00242.XBKR 0000 188
 - space partition 234
 - space XBKR0000 63
 - XBK_TO_STMTR_STMT 64
 - XML node id 57
- XML table space
 - BUFFERPOOL property 231
 - empty or unformatted data pages 191
 - following keywords 219
 - full image copy 202
 - histogram statistics 219
 - inline copies 211
 - NODEID index 253
 - status RW 238
 - table space 259

TRECOVERY option 217
XML TABLESPACE 215
XML type 31, 74–75, 95, 159
 table definition 76
XML type modifier 34, 74, 139, 163, 183, 185, 237, 248
XML value 34, 52, 54, 56, 78, 82, 110, 132–133,
195–196, 198, 220–221, 231, 239, 259
 old versions 59
XML version
 improved storage usage 59
xml version 5, 8–9, 32, 36, 56, 58, 67, 119–120, 122,
157, 197, 199, 221–222, 226, 264, 269
xmladdress c 26
XMLADDRESS table 30
 second row 33
 single XML document 30
XMLAGG 110
XMLELEMENT 25, 109–110, 112, 176
XMLEXISTS 25–26, 108–109, 172, 252–253
XMLEXISTS predicate 27, 109, 254
XMLMODIFY function 32, 81, 124, 167
XMLNAMESPACES 92, 96, 102, 165, 172, 254–255,
260
xmlns 5, 8, 36, 67, 80–81, 119, 190, 264
XMLPARSE 30–32, 59, 84, 115, 159, 177
XMLPARSE function 32
xmlpattern 35, 70, 106, 173, 235, 251, 253
XMLQUERY function 29
XMLR4.SORT WK01 186
XMLR4.SORT WK02 186
XMLR4.SORT WK03 186
XMLR4.SORT WK04 187
XMLSERIALIZE 30–31, 117, 159, 190, 208
XMLTABLE 25, 28, 91–93, 165, 171, 231, 252, 254, 276
XMLTABLE function 28–30, 32, 91, 100, 102–103, 111,
130, 165, 171–172, 178, 254–255, 260
 result table 111
 row XPath expressions 111
XMLXSROBJECTID scalar function 84
XPATH 14–15, 25, 105, 236, 261
XPath 11, 13–14, 22, 24–27, 29–30, 82, 107–108, 165,
169, 172, 231, 236, 244, 246, 251, 259, 261
XPath 2.0 13
XPath expression 18, 27–29, 43, 107–109, 111, 171, 252
 good understanding 35
XPath expressions 13, 29, 35, 111, 252, 261
XPath location 34
 last sequential node 34
XPath patterns 14
XQuery 15, 24–25, 103, 112–115
XQuery expression 15
xs
 element name 10
XSD 6, 9, 11, 37, 44, 67, 77, 89, 138, 162, 179, 190, 235,
251, 264, 276
xsi 67, 80, 190, 208, 255, 264
XSL 11, 276
xsl
 value-of select 18
XSL transformation 12

result tree 12
XSLT 11–13, 50, 129, 137, 149, 276
XSLT processor 13
 desired behavior 13
XSLT stylesheet 18
XSLT transformation 17
XSR 38, 78, 84, 89, 166, 230, 247
XSR object 39

Z

z/OS xix, 4, 19, 89, 130, 163, 192, 225–226, 230, 235,
237, 256–257

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(->Hide:)->Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine:fm still open and **File>Import>Formats the Conditional Text Settings (ONLY!)** to the book files.

Draft Document for Review January 9, 2011 1:25 pm

7915spine.fm 303



Extremely pureXML in DB2 10 for z/OS

(0.5" spine)
0.475" <-> 0.873"
250 <-> 459 pages

To determine the spine width of a book, you divide the paper PPI into the number of pages in the book. An example is a 250 page book using Plainfield opaque 50# smooth which has a PPI of 526. Divided 250 by 526 which equals a spine width of .4752". In this case, you would use the .5" spine. Now select the Spine width for the book and hide the others: **Special>Conditional Text>Show/Hide>SpineSize(->Hide:)>Set** . Move the changed Conditional text settings to all files in your book by opening the book file with the spine:fm still open and **File>Import>Formats the Conditional Text Settings (ONLY)** to the book files.



Extremely pureXML in DB2 10 for z/OS



**Develop Java and
COBOL applications
accessing SQL and
XML data**

**Administer your XML
and SQL data**

**Choose the best
options for
installation and use**

The DB2 pureXML feature offers sophisticated capabilities to store, process and manage XML data in its native hierarchical format. By integrating XML data intact into a relational database structure, users can take full advantage of DB2's relational data management features.

In this IBM Redbooks publication we document the steps for the implementation of a simple but meaningful XML application scenario. We have chosen to provide samples in COBOL and Java language. The purpose being to provide an easy path to follow to integrate the XML data type for the traditional DB2 user.

We have also added considerations for the data administrator and suggested best practices for ease of use and better performance.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:
ibm.com/redbooks**