

An Oracle White Paper
June 2013

Oracle XML DB: Best Practices to Get Optimal Performance out of XML Queries

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Table of Contents

SQL/XML & XQuery	7
Migrating from Oracle Proprietary (XPath 1.0 based) syntax to Standard SQL/XML XQuery based syntax	10
Getting the best performance out of XQuery	14
Storage independent Best Practices.....	15
XQuery Guideline 1: Use XMLExists() and XMLQuery() to search and transform XML stored in XML DB.....	15
XQuery Guideline 2: Use XMLExists() to search the XML document to modify via XML DML operators	15
XQuery Guideline 3: Use XMLTable construct to query XML with relational access	16
XQuery Guideline 4: Use XMLCast() and XMLTable() constructs for GROUP BY and ORDER BY.....	17
XQuery Guideline 5: Use XQuery extension expression to indicate functional evaluation of XQuery.....	18
XQuery Guideline 6: Use XQuery in PL/SQL to manipulate PL/SQL XMLType Variable	18
XQuery Guideline 7: Use proper XQuery and SQL Typing.....	19
XQuery Guideline 8: XQuery expressions that are not optimizable with XML index	22
XQuery Guideline 9: Use the right XQuery expression to access data within Top XQuery	22
XQuery Guideline 10: Gather statistics	25

XQuery Guideline 11: Use SET XMLOPT[IMIZATIONCHECK] or events to determine why a query/DML is not rewritten	26
XQuery Guideline 12: Properly release resources for xmltype in client program	27
XQuery Guideline 13: Avoid calling getObject mutiple times for xmltype in client program	28
XQuery Guideline 14: Set parameter OPTIMIZER_FEATURE_ENABLE to 11.1.0.6 or above for XQuery optimizations	28
Storage dependent performance tuning.....	29
Binary XML.....	30
Binary XML Streaming Evaluation.....	30
Streaming Evaluation Guideline 1: Convert reverse XPath axes to forward axes when possible	31
Streaming Evaluation Guideline 2: For large documents, avoid descendant axis & wild cards if exact (named) path steps can be used	32
Streaming Evaluation Guideline 3: For DML-heavy workloads, enable caching for writes on the underlying lob column	32
Indexing binary XML.....	32
Index choosing Guideline 1: Use the Structured XMLIndex when XPaths are static, and to answer predicates	34
Index choosing Guideline 2: Use Unstructured XMLIndex when XPaths are not known in advance	34
Index choosing Guideline 3: Use Full Text Index for full text search requirements	35

Index choosing Guideline 4: Fragment extraction	35
Index choosing Guideline 5: Combine different indexes as needed ...	36
XMLIndex Structured Component.....	37
Structured Index Guideline 1: Use Structured Index instead of multiple functional indexes and/or virtual columns	38
Structured Index Guideline 2: Make Index and Query datatypes correspond	39
Structured Index Guideline 3: Use XMLTable views with corresponding index, e.g BI style queries	39
Structured Index Guideline 4: Create Secondary Indexes, especially for predicates	41
Structured Index Guideline 5: Check the execution plan to see if structured index is used	41
Structured Index Guideline 6: Indexing Master-Detail relationships	42
Structured Index Guideline 7: Split fragment extraction and value search between SELECT and WHERE clause.....	43
Structured Index Guideline 8: For ordering query results, use SQL ORDER BY along with XMLTable	44
XMLIndex Unstructured Component.....	46
Unstructured XMLIndex Guideline 1: Check the Execution Plan to see if the XMLIndex Unstructured Component is used.....	47
Unstructured XMLIndex Guideline 2: When to drop PIKEY index in favor of ORDERKEY & PATHID index	47
Unstructured XMLIndex Guideline 3: How to use path-subsetting -- smaller index means faster queries	48

Unstructured XMLIndex Guideline 4: Using path-subsetting to choose streaming vs index execution.....	48
Unstructured XMLIndex Guideline 5: Using NO_XMLINDEX_REWRITE_IN_SELECT hint.....	48
Unstructured XMLIndex Guideline 6: Creating datatype aware VALUE indexes by making index and query datatypes correspond	49
Unstructured XMLIndex Guideline 7: XPath Expressions not indexed by Path Subsetted XMLIndex	50
Unstructured XMLIndex Guideline 8: Be specific in the XPath (avoid //, /*)	50
Unstructured XMLIndex Guideline 9: Reduce the number of expressions in the from clause (avoid Path Table join with itself)	50
Unstructured XMLIndex Guideline 10: Use of an index on sys_orderkey_depth	51
Unstructured XMLIndex Guideline 11: Old snapshot queries might be slow	51
Unstructured XMLIndex Guideline 12: Avoid the usage of text() in path expression	51
Full Text Index.....	52
Full Text Index Guideline 1: What To Do If an XML Full-Text Index Is Not Picked Up	54
Full Text Index Guideline 2: Pragma ora:no_schema: Using XML Schema-Based Data with XQuery Full Text	55
Full Text Index Guideline 3: Pragma ora:use_xmltext_idx: Forcing the Use of an XML Full-Text Index	57
Structured (Object Relational) storage.....	59

Structured Storage Guideline 1: Make the SQL types and XML types correspond	59
Structured Storage Guideline 2: Look for underlying tables and columns versus XML functions in execution plans	59
Structured Storage Guideline 3: Name default tables and nested tables, so you recognize them in execution plans	61
Structured Storage Guideline 4: Create relevant indexes	61
Structured Storage Indexing Guideline 5: Create an index on a column targeted by a predicate	61
Structured Storage Indexing Guideline 6: Create indexes on ordered collection tables	64
Querying resources in XML DB repository	67
Appendix A: Semantic differences between the deprecated mainly XPath 1.0 based functions and standard SQL/XML XQuery based functions...	68

Introduction

Oracle XML DB support for the XQuery language is provided through native implementation of SQL/XML functions XMLQuery, XMLTable, XMLEExists, and XMLCast. A SQL statement that includes XMLQuery, XMLTable, XMLEExists, or XMLCast is compiled and optimized as a whole, leveraging both relational database and XQuery-specific optimization technologies.

The XQuery optimizations can be divided into 2 broad areas:

- **Logical optimizations** are transformation of the XQuery into equivalent SQL query blocks extended with XML operators modeling XQuery semantics. These optimizations are generic XQuery optimizations that are independent of the XML storage or indexing model .
- **Physical optimizations** are transformation of the XML operators, in particular, XPath operators, into equivalent operations directly on the underlying internal storage and index tables that are specific to the XML storage and indexing model. The result of XQuery optimization can be examined via explain plan of the SQL/XML query statement that invokes XQuery.

This paper talks about the XQuery Best Practices to get the best performance. It talks about both logical and physical query optimizations. It delves deep into various XML storage and indexing options, and talks about how to choose the right indexes for your query, and how to get the best performance out of your XQuery.

Note: It is assumed that the reader of this whitepaper has already read the whitepaper titled: “Choosing the Best XMLType Storage Option for your Usecase”, and has made the right storage choice for his XMLType data.

SQL/XML & XQuery

Oracle XML DB supports the latest version of the XQuery language specification, i.e., the W3C XQuery 1.0 Recommendation. XQuery 1.0 is the W3C language designed for querying XML data. It is similar to SQL in many ways, but just as SQL is designed for querying structured, relational data, XQuery is designed especially for querying semi-structured, XML data from a variety of data sources. You can use XQuery to query XML data wherever it is found, whether it is stored in database tables, available through Web Services, or otherwise created on the fly. For more information on XQuery 1.0, please see <http://www.w3.org/TR/xquery/>

In addition to XQuery language from W3C, SQL standard has defined standard SQL/XML functions `XMLQuery()`, `XMLExists()`, `XMLCast()` and table construct `XMLTable()` as a general interface between the SQL and XQuery languages. As is the case for the other SQL/XML functions, such as `XMLElement()`, `XMLAgg()`, `XMLForest()`, `XMLConcat()`, that are used to generate XML from relational data, `XMLQuery()`, `XMLExists()`, `XMLCast()` functions and `XMLTable()` table construct let you take advantage of the power and flexibility of both SQL and XML. Using these functions, you can query and manipulate XML, construct XML data using relational data, query relational data as if it were XML, and construct relational data from XML data.

Although SQL/XML functions `XMLQuery()`, `XMLExists()`, `XMLCast()` and `XMLTable()` construct all evaluate an XQuery expression over XMLType input, the way the result of XQuery is consumed varies among them. Therefore, they should be used in the different clauses of SQL to achieve the best performance. In the XQuery language, an expression always returns a sequence of items. The way the sequence of items is consumed in different SQL contexts is classified as below, with the proper usage of these SQL/XML functions and `XMLTable` table construct.

- To consume all the items in the result sequence as a single XML document or fragment, **XMLQuery()** is used as a functional expression, typically in the select list of SELECT clause of SQL, to aggregate the result sequence as one XMLType value representing an XML document or fragment. For example, the query below passes an XMLType column, `oe.warehouse_spec`, as context item to XQuery, using function `XMLQuery` with the `PASSING` clause. It constructs a `Details` element for each of the warehouses whose area is greater than 80,000: `/Warehouse/ Area > 80000`.

Example 1: Using XMLQuery with PASSING Clause

```
SELECT warehouse_name,
       XMLQuery (
         'for $i in /Warehouse
         where $i/Area > 80000
```

```

return <Details>
    <Docks num="{ $i/Docks}"/>
    <Rail>{if ( $i/RailAccess = "Y")
        then "true" else "false"}
    </Rail>
</Details>'
PASSING warehouse_spec RETURNING CONTENT) big_warehouses
FROM warehouses;

```

- **XMLTable()** construct is used in the FROM clause of SQL to return evaluation result of XQuery as a table of rows, each of the XQuery item in the result sequence as an XMLType value. Users can generate a relational view over XML data using XMLTable. This is illustrated below:

Example 2: Using XMLTable to generate a relational view over XML data.

```

SELECT lines.lineitem, lines.description, lines.partid,
       lines.unitprice, lines.quantity
FROM purchaseorder,
     XMLTable('for $i in /PurchaseOrder/LineItems/LineItem
              where $i/@ItemNumber >= 8
              and $i/Part/@UnitPrice > 50
              and $i/Part/@Quantity > 2
              return $i'
              PASSING OBJECT_VALUE
              COLUMNS lineitem      NUMBER      PATH '@ItemNumber',
                       description  VARCHAR2(30) PATH 'Description',
                       partid       NUMBER      PATH 'Part/@Id',
                       unitprice    NUMBER      PATH 'Part/@UnitPrice',
                       quantity     NUMBER      PATH 'Part/@Quantity')
lines;

```

- To determine if XQuery results in empty sequence or not, **XMlexists()**, which has a Boolean result, is typically used in the WHERE or HAVING clause of SQL or conditional expression of SQL CASE expression. The example below shows how XMlexists() can be used in the select list.

Example 3: Using XMlexists() with CASE Expression in select list

```

SELECT
    CASE WHEN XMLEXISTS('$po/PurchaseOrder/LineItems/Part'
    PASSING OBJECT_VALUE AS "po") THEN 1 ELSE 0 END
FROM purchaseorder,

```

- To cast sequence result, typically the leaf value of an XML node, as a SQL scalar type, such as NUMBER, VARCHAR, DATE, TIMESTAMP etc, **XMLCast()** is used as a functional expression resulting in a SQL scalar value item that is used in select list of SELECT clause, group by list of GROUP BY clause, or order by list of ORDER BY clause.

When XQuery is used in SQL/XML functions and XMLTable construct to query XMLType value from tables or views, Oracle XML DB compiles the XQuery expressions into a set of SQL query blocks and operators, and optimizes them by leveraging the underlying XML storage and indexes. This native XQuery/SQL/XML optimization model is achieved conceptually by using a 2-step process: logical optimization and physical optimization.

1. **Logical optimizations** are independent of the XML storage or indexing over the underlying XMLType value. The XQuery expressions that are passed as arguments to SQL/XML functions XMLQuery, XMLExists, XMLCast and XMLTable construct are compiled into internal SQL query blocks and operator trees that model the semantics of XQuery. One common internal operator is the XPath operator that navigates the input XMLType value. A SQL statement that includes XMLQuery, XMLTable, XMLExists, or XMLCast is compiled and optimized as a whole, leveraging both relational database and XQuery-specific optimization technologies.
2. **Physical optimizations** are specific to the underlying storage and indexing model. Depending on the XML storage and indexing methods used, the XPath internal operators can be further optimized into SQL query blocks operating on the underlying physical relational storage tables that are used for the underlying XML index or storage. The relational optimizer optimizes the resulting SQL query blocks and operator trees, in order to achieve the best execution plan.

The resulting query plan is then executed using the SQL row source iterator model. This native XQuery/SQL/XML optimization model achieves the performance goal of primarily using XQuery as a query language to search XML documents stored in the database with the proper XML storage and indexing model, or to present XML as relational results using XMLTable construct. Just as tuning a SQL query using ‘explain plan’ is important, understanding and tuning SQL/XML query using ‘explain plan’ is equally important. This is detailed in the subsequent sections of this document with different XML storage and index options.

Furthermore, XQuery can also be primarily used as a language to manipulate and transform XML documents. The input XMLType value is typically a single XML document or fragment retrieved from persistent XML or transient XMLType value. In such case, XQuery can be functionally evaluated in XML DB. Understanding and classifying XQuery usage in XML DB is critical to get the optimal performance. This is detailed later in this document in the section ‘Getting the best performance out of XQuery.’

Migrating from Oracle Proprietary (XPath 1.0 based) syntax to Standard SQL/XML XQuery based syntax

Starting 11gR2, Oracle has deprecated many older proprietary mainly XPath 1.0 based operators in favor of standards based XQuery syntax, as listed in Table 1 below. If you don't have any code with the functions or operators being deprecated, you may jump to the next section.

TABLE 1. MIGRATING FROM OLD TO XQUERY SYNTAX

OLD ORACLE PROPRIETARY SYNTAX	NEW XQUERY SQL/XML BASED SYNTAX
extract()	XMLQuery()
extractValue	XMLCast(XMLQuery())
existsNode()	XMLExists()
Table (XMLSequence)	XMLTable
ora:instanceof	instanceof
ora:instanceof-only	@xsi:type
getNamespace	fn:namespace-uri
getRootElement	fn:local-name
getStringVal, getBlobVal, getClobVal	XMLSerialize
Xmltype()	XMLParse() for varchar, clob, blob input
DBMS_XMLQUERY	XMLQuery()
DBMS_XMLGEN	SQL/XML Operators
Oracle XML DML Operators	XQuery Update Facility

There are some important semantic differences between the deprecated mainly XPath 1.0 based syntax and the XQuery based syntax. These are listed in Appendix A to make the migration easier for the users. Please also check “XQuery Guideline 6” in this document to see how to apply XQuery to PL/SQL XMLType variable instead of calling extract() and existsNode() methods of xmltype.

The table below shows examples of Oracle Proprietary XML DML operators and their equivalent Xquery Update syntax:

Note: Oracle Proprietary XMLDML does not have “rename” and “insert as first into” operations.

TABLE 2. MIGRATING FROM ORALCE XMLDML TO XQUERY UPDATE SYNTAX

ORIGINAL EXAMPLE	REPLACEMENT EXAMPLE
Update warehouses set warehouse_spec = appendChildXML (warehouse_spec, 'Warehouse/Building', XMLType('<Owner>Grandco</Owner>'));	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify (for \$i in \$tmp/Warehouse/Building return insert node <Owner>Grandco</Owner> as last into \$i) return \$tmp' passing warehouse_spec returning content);
Update warehouses set warehouse_spec = deleteXML (value(po), '/Warehouse/Building');	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify delete node \$tmp/Warehouse/Building return \$tmp' passing warehouse_spec returning content);
[Single Node Case] Update warehouses set warehouse_spec = insertXML (warehouse_spec, '/Warehouse/Building/Owner[2]', XMLType('<Owner>ThirdOwner</Owner>'));	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify insert node <Owner>ThirdOwner</Owner> into \$tmp/Warehouse/Building/Owner[2] return \$tmp' passing warehouse_spec returning content);
[Single Node Case] Update warehouses set warehouse_spec = insertXMLBefore (warehouse_spec, '/Warehouse/Building/Owner[2]', XMLType('<Owner>FirstOwner</Owner>'));	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify insert node <Owner>FirstOwner</Owner> before \$tmp/Warehouse/Building/Owner[2] return \$tmp' passing warehouse_spec returning content);
[Single Node Case] Update warehouses set warehouse_spec = insertXMLAfter (warehouse_spec, '/Warehouse/Building/Owner[2]', XMLType('<Owner>ThirdOwner</Owner>'));	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify insert node <Owner>ThirdOwner</Owner> after \$tmp/Warehouse/Building/Owner[2] return \$tmp' passing warehouse_spec returning content);
Update warehouses set warehouse_spec = updateXML (warehouse_spec, '/Warehouse/Docks/text()', 4);	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify (for \$i in \$tmp/Warehouse/Docks/text() return replace value of node \$i with 4) return \$tmp' passing warehouse_spec returning content);
Update warehouses set warehouse_spec = insertChildXML (warehouse_spec, '/Warehouse/Building', 'Owner', XMLType('<Owner>LesserCo</Owner>'));	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify (for \$i in \$tmp/Warehouse/Building return insert node <Owner>LesserCo</Owner> into \$i) return \$tmp' passing warehouse_spec returning content);
Update warehouses set warehouse_spec = insertChildXMLBefore (warehouse_spec, '/Warehouse/Building', 'Owner',	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify (for \$i in \$tmp/Warehouse/Building return insert node

XMLType('<Owner>LesserCo</Owner>');	<Owner>LesserCo</Owner> before \$i) return \$tmp' passing warehouse_spec returning content);
Update warehouses set warehouse_spec = insertChildXMLAfter (warehouse_spec, '/Warehouse/Building', 'Owner', XMLType('<Owner>LesserCo</Owner>'));	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify (for \$i in \$tmp/Warehouse/Building return insert node <Owner>LesserCo</Owner> after \$i) return \$tmp' passing warehouse_spec returning content);
[Collection Case] Update warehouses set warehouse_spec = insertXML (warehouse_spec, '/Warehouse/Building/Owner', XMLType('<Owner>AnotherOwner</Owner>'));	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify (for \$i in \$tmp/Warehouse/Building/Owner return insert node <Owner>AnotherOwner</Owner> into \$i) return \$tmp' passing warehouse_spec returning content);
[NULL Case] Update warehouses set warehouse_spec = updateXML (warehouse_spec, '/Warehouse/Docks', null);	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify delete node \$tmp/Warehouse/Docks return \$tmp' passing warehouse_spec returning content);
[Empty Node Case] Update warehouses set warehouse_spec = updateXML (warehouse_spec, '/Warehouse/Docks', '');	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := \$p1 modify (for \$j in \$tmp/Warehouse/Docks return replace node \$j with \$p2) return \$i' passing passing warehouse_spec "p1", " as "p2" returning content) ;
[Multiple Path Case] Update warehouses set warehouse_spec = updateXML (warehouse_spec, '/Warehouse/Docks/text()', extractValue(warehouse_spec, '/Warehouse/Docks/text()')+4, '/Warehouse/Docks/text()', extractValue(warehouse_spec, '/Warehouse/Docks/text()')+4);	Update warehouses set warehouse_spec = XMLQuery('copy \$tmp := . modify ((for \$i in \$tmp/Warehouse/Docks/text() return replace value of node \$i with \$i+4), (for \$i in \$tmp/Warehouse/Docks/text() return replace value of node \$i with \$i+4)) return \$tmp' passing warehouse_spec returning content);

The XQuery and XPath Full Text (XQFT) standard is supported by Oracle XML DB starting with Oracle Database 12c Release 1 (12.1). Prior to this release, for full-text querying of XML data you could use only an Oracle Text index that was not XML-enabled, and your full-text queries necessarily used Oracle-specific constructs: SQL function CONTAINS or XPath function ora:contains.

If you have legacy code that does this, Oracle recommends that you migrate that code to use XQFT. The table below provides a mapping from typical queries that use Oracle-specific constructs to queries that use XQuery Full Text.

TABLE 3. MIGRATING FROM ORACLE CONTAINS TO XQUERY FULL TEXT SYNTAX

ORIGINAL EXAMPLE	REPLACEMENT EXAMPLE
CONTAINS(t.x, 'HASPATH (/P/LIs/LI/Description)') > 0	XMLExists('(/# ora:use_xmltext_idx #) {\$d/P/LIs/LI/Description}')

ORIGINAL EXAMPLE	REPLACEMENT EXAMPLE
	<p>PASSING t.x AS "d")</p> <p>Or if the data is XML Schema-based:</p> <pre>XMLExists('(# ora:use_xmltext_idx #) {(# ora:no_schema #) {\$d/P/Ls/LI/Description})' PASSING t.x AS "d")</pre>
<pre>CONTAINS(t.x, 'Big INPATH (/P/Ls/LI/Description)') > 0</pre>	<pre>XMLExists('\$d/P/Ls/LI/Description [. contains text "Big"] PASSING t.x AS "d")</pre> <p>Or if the data is XML Schema-based:</p> <pre>XMLExists('(# ora:no_schema #) {\$d/P/Ls/LI/Description [. contains text "Big"])' PASSING t.x AS "d")</pre>
<pre>XMLExists('\$d/P/Ls/LI/Description [ora:contains(., "Big AND Street") > 0] PASSING t.x AS "d")</pre>	<pre>XMLExists('\$d/P/Ls/LI/Description [. contains text "Big" ftand "Street"] PASSING t.x AS "d")</pre>
<pre>CONTAINS(t.x, '(Big) AND (Street) INPATH (/P/Ls/LI/Description)') > 0</pre>	<pre>XMLExists('\$d/P/Ls/LI/Description [. contains text "Big" ftand "Street"] PASSING t.x AS "d")</pre>
<pre>CONTAINS(t.x, '(Big) OR (Street) INPATH (/P/Ls/LI/Description)') > 0</pre>	<pre>XMLExists('\$d/P/Ls/LI/Description [. contains text "Big" ftor "Street"] PASSING t.x AS "d")</pre>
<pre>CONTAINS(t.x, '({Big}) NOT ({Street}) INPATH (/P/Ls/LI/Description)') > 0</pre>	<pre>XMLExists('\$d/P/Ls/LI/Description [. contains text "Big" ftand ftnot "Street"] PASSING t.x AS "d")</pre>
<pre>CONTAINS(t.x, '({Street}) MNOT ({Big Street}) INPATH (/P/Ls/LI/Description)') > 0</pre>	<pre>XMLExists('\$d/P/Ls/LI/Description [. contains text "Street" not in "Big Street"] PASSING t.x AS "d")</pre>
<pre>CONTAINS(t.x, '(NEAR (({Big}, {Street}), 3) INPATH</pre>	<pre>XMLExists('\$d/P/Ls/LI/Description</pre>

ORIGINAL EXAMPLE	REPLACEMENT EXAMPLE
<code>(/P/LIs/LI/Description)) > 0</code>	<pre>[. contains text "Big" fband "Street" window 3 words] PASSING t.x AS "d")</pre>
(Not applicable – Oracle Text queries are not XML namespace aware.)	<pre>XMLExists('declare namespace ipo="http://www.example.com/IPO"; /ipo:P/ipo:LIs/ipo:LI/ipo:Description [. contains text "Big"] PASSING t.x AS "d")</pre>

Getting the best performance out of XQuery

XQuery Best Practices and Performance Tuning can be divided into 2 parts:

- Best practices independent of the XMLType storage options. These are listed in the “Storage independent best practices” section.
- Best practices specific to the XMLType storage selected by the user. These include various indexes the user can create to speed up their XQueries. These are listed in the “Storage dependent performance tuning” section.

Storage independent Best Practices

In Oracle XML DB, XML documents are stored in either XMLType tables or XMLType columns of relational tables. XML DB is designed to store large number of XML documents, and to search using XQuery among these XML documents, in order to find qualified XML documents or document fragments for manipulation and transformation using XQuery, or to project relational views over XML using XMLTable construct so that they can be queried relationally and be integrated with mature relational applications.

XQuery Guideline 1: Use XMLEExists() and XMLQuery() to search and transform XML stored in XML DB

The typical way of writing a SQL statement that searches XML documents stored in XMLType column and manipulates the searched result is stated below:

Example 4: Search and transform

```
SELECT XMLQUERY ('...' PASSING T.X RETURNING CONTENT)
FROM purchaseorder T
WHERE XMLEXISTS ('$p/PurchaseOrder/LineItems/LineItem/Part[@Id="702372"]'
                PASSING T.X AS "p");
```

In this SQL statement, XMLEExists() is used in WHERE clause of the statement to accomplish the typical database task of “finding needle in a haystack.” Since there can be billions of XML document stored in table purchaseorder, using proper index, instead of a table scan with functional evaluation of XQuery used in XMLEExists() for each XML document, is critical to achieve query performance. To achieve the best performance, the XQuery used in XMLEExists() should be index friendly so that when the XMLType column is stored using structured storage, relational index created on the underlying relational tables from the structured storage is used, and when the XMLType column is stored using binary XML storage, XMLIndex over the XML storage is used.

If XQuery used in XMLEExists() is not index friendly as a whole, then try to break the XQuery into index-friendly expressions and index-unfriendly expressions and use them in two different XMLEExists() functions connected by the SQL AND construct. In this way, at least the index-friendly XMLEExists() can be evaluated using index and the index-unfriendly XMLEExists() can be evaluated as a post-index filter.

XQuery Guideline 2: Use XMLEExists() to search the XML document to modify via XML DML operators

The typical way of writing a SQL statement that searches for and modifies XML documents stored in XMLType column is shown below.

Example 5: Updating XML document after searching using XMLEExists()

```

UPDATE purchaseorder T SET T.X = DELETXML(T.X,
      '/purchaseOrder/LineItems/LineItem[itemName = "TV"]' )
WHERE
XML EXISTS ('$p/PurchaseOrder/LineItems/LineItem/Part[@Id="717951002372"]'
      PASSING T.X AS "p");

```

As in XQuery Guideline 1, XMLExists() is used here to identify which XML documents is to be modified, i.e., “finding needle in a haystack.” The function used in the RHS of the UPDATE assignment can be any expression that returns XMLType. For example, it can be a PL/SQL function call that returns XMLType. Semantically, the RHS expression of the SQL UPDATE statement returns an XMLType instance document that is assigned to XMLType column on the LHS to do document replacement of the whole XMLType column value.

However, Oracle XML DB does XML DML operator rewrite optimization whenever possible, so as to partially update the underlying XML storage structures instead of replacing the whole document, as listed below:

- For binary XML storage, there is XML DML operator rewrite for all XML DML operators when the XPath can be evaluated using streaming evaluation.
- For structured storage, there is XML DML operator rewrite for DELETXML(), UPDATEXML(), INSERTCHILDXML(), INSERTCHILDXMLBEFORE(), and INSERTCHILDXMLAFTER() ,when the XPath can be evaluated directly using the underlying relational tables and columns of the structured storage.

XML DML operator rewrite can be explicitly disabled by using the /*+NO_XML_DML_REWRITE*/ SQL hint. This is true regardless of XML storage model.

XQuery Guideline 3: Use XMLTable construct to query XML with relational access

XML document is hierarchical in nature and has typical master-detail relationships. Therefore, it is common to project out master-detail constructs within XML document as a set of relational tables using XMLTable construct and project out leaf values of each construct as columns of XMLTable for search, as shown in the example below:

Example 6: Using XMLTable

```

SELECT li.description, li.lineitemFROM purchaseorder T,
      XMLTable('$p/PurchaseOrder/LineItems/LineItem'
      PASSING T.X AS "p"
      COLUMNS lineitem    NUMBER           PATH '@ItemNumber',
      description  VARCHAR2(30)  PATH 'Description',
      partid       NUMBER         PATH 'Part/@Id',
      unitprice    NUMBER         PATH 'Part/@UnitPrice',
      quantity     NUMBER         PATH 'Part/@Quantity') li
WHERE li.unitprice > 30 and li.quantity < 20);

```

To process the XMLTable() construct efficiently, XQuery usage in XMLTable clause should be storage or index friendly so that native XQuery/SQL/XML optimization can find the best query plan leveraging the underlying XML storage and index models. In this case, if purchaseorder column is stored using structured storage, the underlying relational table holding the LineItem together with its relational columns unitprice and quantity are directly accessed in the resulting query plan. If purchaseorder column is stored using binary XML storage, the underlying relational tables belonging to the XMLIndex are directly accessed in the resulting query plan.

To traverse multi-level hierarchy, XMLTable can be used in a chaining fashion.

XQuery Guideline 4: Use XMLCast() and XMLTable() constructs for GROUP BY and ORDER BY

There are GROUP BY and ORDER BY clauses that operate on SQL scalar types. One typical way of casting XQuery result into SQL scalar types for GROUP BY and ORDER BY purposes is shown in the example below.

Example 7: Using XMLCast() in GROUP BY / ORDER BY

```
SELECT XMLCAST(XMLQUERY('$p/PurchaseOrder/@poDate' PASSING T.X
                      RETURNING CONTENT) AS DATE), COUNT(*)
FROM purchaseorder T
WHERE ...
GROUP BY XMLCAST(XMLQUERY('$p/PurchaseOrder/@poDate' PASSING T.X
                      RETURNING CONTENT) AS DATE)
ORDER BY XMLCAST(XMLQUERY('$p/PurchaseOrder/@poDate' PASSING T.X
                      RETURNING CONTENT) AS DATE);
```

When there are multiple scalar values that need to be grouped or ordered, it is better to write it with XMLTable construct that projects out all columns to be ordered or grouped as shown below.

Example 8: Using XMLTable() construct for GROUP BY / ORDER BY

```
SELECT po.DATE, po.poZip, count(*)
FROM purchaseorder T,
     XMLTable('$p/PurchaseOrder'
              PASSING T.X AS "p"
              COLUMNS
                 poDate      DATE          PATH '@poDate',
                 poZip       VARCHAR2(8)   PATH 'shipAddress/zipCode',
              ) po
WHERE ...
GROUP BY po.poDate, po.poZip
ORDER BY po.poDate, po.poZip
```

In this case, if purchaseOrder.X column uses structured storage or uses binary XML storage with structured xmlindex, the query plan will directly use group by and order by of the columns from the underlying relational storage tables of the XML storage or xmlindex.

Note the XMLTable usage pattern in SQL/XML is very commonly adopted by users to create relational views over XML, so that XML query can be integrated with existing relational applications (such as BI applications) smoothly.

XQuery Guideline 5: Use XQuery extension expression to indicate functional evaluation of XQuery

XQuery is a language that blends both search and transformation of XML. While XQuery used for search in the WHERE clause is more amenable for XQuery rewrite optimization leveraging the underlying XML storage and indexing models, XQuery used for transformation in the SELECT clause might be more procedure-centric and hence suited for functional evaluation. You can use XQuery extension expression (`#ora:xq_proc #`) to indicate that the XQuery should be functionally evaluated, as shown in the example below.

Example 9: XQuery extension expression for functional evaluation

```
SELECT XMLQUERY('#ora:xq_proc #}{...}' PASSING T.X RETURNING CONTENT)
FROM purchaseorder T
WHERE XMLEXISTS('$p/PurchaseOrder/LineItems/LineItem/Part[@Id="717951"]'
                PASSING T.X AS "p");
```

The `(#ora:xq_proc#){...}` is an XQuery extension expression serving as a “pragma” to indicate the xquery expression enclosed in the curly braces needs to be evaluated functionally. It is available since Oracle 11gR2, release 11.2.0.2.

This mechanism is more fine-grained and hence more flexible than using `/*+ NO_XML_QUERY_REWRITE */` SQL hint, which requests all XQuery used in a SQL statement to use functional evaluation. This may not be desirable for XQuery used in `XMLEXISTS()` of the SQL statement.

XQuery Guideline 6: Use XQuery in PL/SQL to manipulate PL/SQL XMLType Variable

PL/SQL XMLType methods do not support XQuery invocation directly. However, one can invoke SQL/XML functions with XQuery to query on XMLType PL/SQL variables as shown in the following example. Since a PL/SQL XMLType variable value is not indexed, `/*+ NO_XML_QUERY_REWRITE*/` SQL hint is used to evaluate XQuery functionally.

Example 10: Querying PL/SQL XMLType variable using XMLQuery() and XMLCast()

```
DECLARE
    v_x XMLType;
    NumAcc NUMBER;
BEGIN
    v_x := XMLType(xmlfile(...)); /* initialize xmltype variable */
    SELECT /*+ NO_XML_QUERY_REWRITE */
           XMLCAST(XMLQUERY('declare default element namespace
```

```

        "http://custacc";for $cust in $cadoc/Customer return
        fn:count($cust/Addresses/Address) '
    PASSING v_x AS "cadoc" RETURNING CONTENT) AS NUMBER)
    INTO NumAcc
    FROM DUAL;
END;

```

Example 11: Querying PL/SQL XMLType variable using XMLEExists()

```

DECLARE
    v_x XMLType;
    ex NUMBER;
BEGIN
    v_x := XMLType(xmlfile(...)); /* initialize xmltype variable */
    SELECT /*+ NO_XML_QUERY_REWRITE */
        CASE WHEN XMLEXISTS('declare default element namespace
            "http://custacc"; $cadoc/Customer/Addresses/Address) '
            PASSING v_x AS "cadoc")
            THEN 1 ELSE 0 END
    INTO ex
    FROM DUAL;
END;

```

XQuery Guideline 7: Use proper XQuery and SQL Typing

XQuery type system is based on XML Schema type system. Although XQuery type system and SQL type system are not exactly aligned, there are equivalent mappings between the types in each system, as shown in the table below. Note that xs:date, xs:time, xs:dateTime have optional timezone component, therefore, they are mapped to 'TIMESTAMP WITH TIMEZONE' SQL type. When the timezone component is not used, then you may map to DATE or TIMESTAMP SQL types.

TABLE 2. XML AND SQL DATA TYPE CORRESPONDENCE FOR XMLINDEX

XML DATA TYPE	SQL DATA TYPE
xs:integer, xs:decimal	INTEGER or NUMBER
xs:double	BINARY_DOUBLE
xs:float	BINARY_FLOAT
xs:date	DATE, TIMESTAMP WITH TIMEZONE

xs:time, xs:dateTime	TIMESTAMP, TIMESTAMP WITH TIMEZONE
xs:dayTimeDuration	INTERVAL DAY TO SECOND
xs:yearMonthDuration	INTERVAL YEAR TO MONTH

Users are recommended to cast these types properly in XQuery used within `XMLExists()` clause to ensure proper type-aware comparison semantics and proper XML index usage. This is illustrated in the examples below:

Example 12: Using XQuery type casting and SQL type cast to pass in the properly typed value into `XMLExists()`

```
SELECT ... FROM purchaseOrder T
WHERE XML EXISTS ('$po/purchaseOrder[@id=$id]'
                 PASSING T.X AS "po", CAST(:1 AS NUMBER) as "id" );
```

In this example, we explicitly cast SQL bind variable `:1` as SQL NUMBER type and bind that to XQuery external variable “`$id`” of `XMLExists()` operator.

If the `purchaseOrder` document is non-XML schema based, then `@id` is of type `xs:untypedAtomic`. The general comparison rule in XQuery states that comparing `xs:untypedAtomic` value with any numeric type value (`xs:integer`, `xs:decimal`, `xs:float`, `xs:double`) is done by promoting both operands to `xs:double`. This makes the `@id` comparison in XQuery use `xs:double()` comparison even though SQL bind variable is passed as `xs:decimal` typed value, it is internally casted into `xs:double` typed value.

On the other hand, if the `purchaseOrder` document is XML schema based, then `@id` is not of type `xs:untypedAtomic`, instead it is of type stated by the XML Schema. If the XML schema states that the `@id` is of type `xs:decimal`, for example, then this makes the `@id` comparison in XQuery use `xs:decimal()` comparison and the SQL bind variable passed as `xs:decimal` typed value no longer needs to be internally casted into `xs:double` typed value.

Keeping in mind that `xs:decimal` is for exact numeric type and `xs:double` is for approximate numeric type, application users need to decide what typed comparison the application needs. Once the decision is made, then write the “Example 12: Using XQuery type casting and SQL type cast to pass in the properly typed value into `XMLExists()`” query above as “Example 13: using `xs:decimal()` type exact numeric comparison” or “Example 14: Using `xs:double()` type approximate numeric comparison” using explicit XQuery type casting to get either `xs:decimal()` typed comparison or `xs:double()` typed comparison.

Example 13: using `xs:decimal()` type exact numeric comparison

```
SELECT ... FROM purchaseOrder T
WHERE XML EXISTS ('$po/purchaseOrder[xs:decimal(@id)=$id]'
                 PASSING T.X AS "po", CAST(:1 AS NUMBER) as "id" );
```

Example 14: Using xs:double() type approximate numeric comparison

```
SELECT ... FROM purchaseOrder T
WHERE XMLEXISTS ( '$po/purchaseOrder[xs:double(@id)=$id]'
                  PASSING T.X AS "po", CAST(:1 AS BINARY_DOUBLE) as
                  "id" );
```

Using explicit type casting is required to ensure that XQuery will use proper typed value comparison independent of whether XMLType document stored in the table is XML schema based or not. Furthermore, doing so promotes the usage of XMLIndex.

To make “Example 13: using xs:decimal() type exact numeric comparison” use structured XMLIndex, “/purchaseOrder/@id” must be indexed as SQL NUMBER type.

To make “Example 13: using xs:decimal() type exact numeric comparison” use unstructured XMLIndex, DBMS_XMLINDEX.CreateNumberIndex() must be called with ‘DECIMAL’ as xmltypename parameter.

To make “Example 14: Using xs:double() type approximate numeric comparison” use structured XMLIndex, “/purchaseOrder/@id” must be indexed as SQL TO_BINARY_DOUBLE type.

To make “Example 14: Using xs:double() type approximate numeric comparison” use unstructured XMLIndex, DBMS_XMLINDEX.CreateNumberIndex() must be called with ‘DOUBLE’ as xmltypename parameter.

See structured and unstructured XMLIndex guideline sections for details.

For non-numeric datatypes, XQuery general comparison allows xs:untypedAtomic typed value to be cast into the type of the other value, so we just need to apply XQuery type casting on the passing parameter as shown in the 2 examples below for xs:date() and xs:dateTime() comparison.

Example 15: Using xs:date() for date datatype comparison

```
SELECT ... FROM purchaseOrder T
WHERE XMLEXISTS ( '$po/purchaseOrder[@podate =xs:date($d)]'
                  PASSING T.X AS "po", :1 as "d" );
```

Here, :1 is expected to bind to SQL varchar of value, say ‘2008-07-08’.

Example 16: Using xs:dateTime() for timestamp with timezone datatype comparison

```
SELECT ... FROM purchaseOrder T
WHERE XMLEXISTS ( '$po/purchaseOrder[@podate =xs:dateTime($d)]'
                  PASSING T.X AS "po", :1 as "d" );
```

Here, :1 is expected to bind to SQL varchar of value, say ‘2010-01-01T12:00:00Z’.

XQuery Guideline 8: XQuery expressions that are not optimizable with XML index

Some expressions might add performance overhead when processing large-size XML document, because these expressions typically cannot leverage the underlying XML storage or index structures. Such expressions should be avoided when querying very large XML documents. They are listed in Table 3:

TABLE 3. EXPRESSIONS TO AVOID FOR LARGE DOCUMENTS

EXPRESSIONS TO AVOID
Avoid XQuery expressions that use the following XPath step axes:
<ul style="list-style-type: none"> • ancestor • ancestor-or-self • descendant-or-self • following • following-sibling • parent • preceding • preceding-sibling
Avoid <<, >> expressions.

XQuery Guideline 9: Use the right XQuery expression to access data within Top XQuery

Pure XQuery users prefer to write XQuery without using individual SQL/XML operators. Oracle XML DB supports this type of usage by enabling users to wrap the entire XQuery into one SQL SELECT statement using either

```
SELECT * FROM XMLTABLE('...');
```

or

```
SELECT XMLQuery('...') FROM DUAL;
```

depending on whether the XQuery results are consumed as a sequence or as one XML fragment. This is referred to as “Top XQuery” because SQL is used here purely as a wrapping mechanism.

Prior to 11gR2 11.2.0.2 release, functions `fn:collection()` and `fn:doc()` needed to be replaced with `ora:view()`. In 11gR2 11.2.0.2 release, `fn:collection()` or `fn:doc()` can be used to uniformly refer to XML documents that are stored in XMLType tables, XMLType columns, XML DB repository, or generated virtually from pure relational tables. However, you need to use the proper `oradb-`prefixed URL or XQuery extension expression. Examples are shown below.

Top XQuery statement goes through the same XQuery rewrite optimizations as that of regular SQL/XML statements. Just as users do performance tuning using `explain plan` for SQL

statements, users should use explain plan to do performance tuning for Top XQuery statement as well.

- Use ora:view() to map relational table content as a collection of virtual XML documents

```
SELECT *
FROM XMLTABLE(
  'for $i in ora:view("SCOTT", "EMP")
  where $i/ROW[EMPNO = 7369 and HIREDATE=xs:date("1980-12-17")]
  return $i' );
```

Here EMP is a relational table owned by user “SCOTT”.

In 11gR2 11.2.0.2 release or later, you may also use fn:collection() as shown below :

```
SELECT * FROM XMLTABLE(
  'for $i in fn:collection("oradb:/SCOTT/EMP")
  where $i/ROW[EMPNO = 7369 and HIREDATE=xs:date("1980-12-17")]
  return $i' );
```

- Use ora:view() to map XMLType table content as a collection of XML documents

```
SELECT * FROM XMLTABLE(
  'for $i in ora:view("PO", "PURCHASEORDER")
  where $i/PurchaseOrder/Id = xs:decimal(789645)
  return $i/PurchaseOrder/LineItems/LineItem[itemName="TV"]' )
```

Here, PURCHASEORDER is an XMLType table owned by user PO.

In 11gR2 11.2.0.2 release, you may also use fn:collection() as shown below:

```
SELECT * FROM XMLTABLE(
  'for $i in fn:collection("oradb:/PO/PURCHASEORDER")
  where $i/PurchaseOrder/Id = xs:decimal(789645)
  return $i/PurchaseOrder/LineItems/LineItem[itemName="TV"]' )
```

- Use fn:collection() to map XMLType column of a table as a collection of XML documents

Here, PURCHASEORDER is a relational table owned by user PO and has an XMLType column ‘X’. This is available starting 11gR2 11.2.0.2 release.

```
SELECT * FROM XMLTABLE(
  'for $i in fn:collection("oradb:/PO/PURCHASEORDER/ROW/X")
  where $i/PurchaseOrder/Id = xs:decimal(789645)
  return $i/PurchaseOrder/LineItems/LineItem[itemName="TV"]' )
```

- Use XQuery extension expression with `fn:collection()` and `fn:doc()` to map XML documents stored in XML DB repository.

When XML documents are stored in XML DB repository (`RESOURCE_VIEW`), they have file/folder paths and can be efficiently accessed using `equals_path()` and `under_path()` SQL functions over `RESOURCE_VIEW`. Prior to 11gR2 release 11.2.0.2, Oracle does not recommend using Top XQuery to access XML documents stored in XML DB repository. However, starting with 11gR2 release 11.2.0.2 release, you may use Top XQuery with proper XQuery extension expression to access them.

In Oracle 11gR2, release 11.2.0.2, Oracle XQuery extension expression `ora:defaultTable` lets you specify the default table used to store repository data that you query. When this XQuery extension expression is used with XQuery functions `fn:doc` and `fn:collection`, the query is rewritten to automatically join the default table to view `RESOURCE_VIEW` and use Oracle SQL functions `equals_path` and `under_path`, respectively. The effect is thus the same as coding the query manually to use an explicit join. Avoid `fn:doc` and `fn:collection` that do not include the `ora:defaultTable` XQuery pragma, since such XQuery will not be optimized.

Example 17: Using `fn:doc` or `fn:collection` with `ora:defaultTable` pragma

```
SELECT XMLQuery (' (#ora:defaultTable PURCHASEORDER #)
{let $val :=fn:doc ("/home/OE/PurchaseOrders/2002/Sep/VJONES-
20021009123337583PDT.xml") /PurchaseOrder/LineItems/LineItem[@ItemN
umber =19] return $val}' RETURNING CONTENT)
FROM DUAL;
```

Example 18: How *not* to write your XQuery with `fn:doc` or `fn:collection`

```
SELECT XMLQuery ('let $val :=
fn:doc ("/home/OE/PurchaseOrders/2002/Sep/VJONES-
20021009123337583PDT.xml")
/PurchaseOrder/LineItems/LineItem[@ItemNumber =19]
return $val' RETURNING CONTENT)
FROM DUAL;
```

In releases prior to Oracle 11gR2, release 11.2.0.2, avoid `fn:doc` and `fn:collection`.

Instead, join the view `RESOURCE_VIEW` with the `XMLType` table that holds the data, and then use the Oracle-specific SQL functions `equals_path` and `under_path` instead of the XQuery functions `fn:doc` and `fn:collection`, respectively. These SQL functions reference repository resources in an efficient way.

Example 19: Accessing documents stored in repository using `RESOURCE_VIEW`

```
SELECT XMLQuery ('let $val :=
$DOC/PurchaseOrder/LineItems/LineItem[@ItemNumber = 19]
return $val' PASSING OBJECT_VALUE AS "DOC" RETURNING CONTENT)
```

```

FROM RESOURCE_VIEW rv, purchaseorder x
WHERE ref(x) = XMLCast(XMLQuery('declare default element namespace
"http://xmlns.oracle.com/xdb/XDBResource.xsd"; (: :) fn:data
(/Resource/XMLRef)' PASSING rv.RES RETURNING CONTENT)
AS REF XMLType)
AND equals_path(rv.RES, '/home/OE/PurchaseOrders/2002/Sep/VJONES-
20021009123337583PDT.xml')

```

- To avoid passing hard-coded search values as constants to Top-XQuery, users may use PASSING bind variable parameters as shown the example below:

Example 20: Passing Bind Variables

```

SELECT * FROM XMLTABLE(
    'for $i in fn:collection("oradb:/SCOTT/EMP")
    where $i/ROW[EMPNO = xs:decimal($empno)]
    return $i'
    PASSING :1 as "empno")

```

XQuery Guideline 10: Gather statistics

One common problem is that user forgets to gather stats on his tables. Inaccurate stats can result in a bad execution plan. Hence it is recommended to periodically perform gather statistics on the XMLType table and relevant indexes, as listed below.

In a use case where data is loaded once and queried several times, running `dbms_stats.gather_table_stats()` on the affected tables (as outlined below), after data has been loaded, is sufficient. In a use case where data is loaded or updated quite frequently, running `dbms_stats.gather_schema_stats()` or `dbms_stats.gather_table_stats()` (as outlined below) as a background scheduler job (package `dbms_scheduler`) is the best. Note that the default behavior of `gather_table_stats` is to propagate gathering of stats to all indexes on the table.:

- For structured storage, user needs to gather stats explicitly on top-level table, each of the nested tables, and each of the out-of-line tables. Gathering stats on top-level table does not result in automatically gathering stats on nested tables and out-of-line tables.
- For XMLIndex, gathering stats on base table will automatically gather stats on the Structured XMLIndex tables and Path table of Unstructured XMLIndex. Hence, there is no need to gather stats on the XMLIndex separately.
- For Text Index, gathering stats on base table will automatically gather stats on the Text Index tables. Hence, there is no need to gather stats on Text Index separately.

Starting Oracle 11.2.0.3, if there are xml indexes present that use binary-double as secondary indexes, it is recommended to set `optimizer_dynamic_sampling` to 3 for picking up proper secondary indexes. For example, the following 2-command script can be used to gather statistics on the schema:

```
alter session set optimizer_dynamic_sampling = 3;
exec dbms_stats.gather_schema_stats('USERNAME');
```

XQuery Guideline 11: Use SET XMLOPT[IMIZATIONCHECK] or events to determine why a query/DML is not rewritten

Just as query tuning can improve SQL performance, so it can improve XQuery performance. You tune XQuery performance by choosing appropriate indexes for your XML Storage. As with database queries generally, you can examine the execution plan for a query to determine whether tuning is required.

In general, use explain plan on your SQL statement (including Top XQuery wrapped in SQL statement) to understand and tune query performance. In particular, when there is 'COLLECTION ITERATOR' appearing in the explain plan, it usually indicate the query plan is not fully optimized.

Advanced users can use:

- XMLOPT[IMIZATIONCHECK] [ON|OFF]" mechanism (in Oracle 11gR2 release 11.2.0.2), or event 19021 with level 4096 (0x1000)(in releases prior to 11.2.0.2) to get the optimized rewritten query in the trace file to see what underlying queries are executed on the underlying internal tables created for XML storage and index models.
- Event 19027 with level 8192 (0x2000) to get a dump in the trace file indicating why a particular expression is not rewritten.

In Oracle 11gR2, release 11.2.0.2, or later:

In Oracle 11gR2 11.2.0.2 release or later, we recommend that you use the "SET XMLOPT[IMIZATIONCHECK] [ON|OFF]" mechanism to determine if parts of your query were not optimized. When it is ON, it will ensure that only XML queries or XML operations that were fully optimized will be executed. A suboptimal XML query or DML operation will be aborted with the following error message: "ORA 19022 - Unoptimized XML construct detected". In addition, the reason for the query or DML being suboptimal will be printed to the trace file. OFF will not guarantee that only fully optimized XML queries/ DML operations will be executed. The default option for this command is OFF. Please use XMLOPT[IMIZATIONCHECK] ON only when developing or debugging a query/ DML operation for performance tuning.

In releases prior to Oracle 11gR2 11.2.0.2 :

If you are on a release prior to 11gR2 11.2.0.2 release, you may set event 19021 with level 1 for a given database session using SQL statement ALTER SESSION to determine if your XML

operation was rewritten. Turn on event 19021 with level 1 if you want to raise an error whenever any of the XML functions is not rewritten and is instead evaluated functionally. The error “ORA-19022 - XML XPath functions are disabled” is raised when such functions execute.

XQuery Guideline 12: Properly release resources for xmltype in client program

When XMLType result is fetched in JDBC program, please make sure to call close() method, or free() method if you are using SQLXML interface, on XMLType result once it is consumed to free resources allocated by the server to track the XMLType results. The following JDBC code fragment demonstrates the call of close() method on XMLType result.

Example 21: Using the close() method to free the resources in JDBC

```
XMLType xml2;
while (rset.next())
{
    xml2 = (XMLType)rset.getObject(1);
    System.out.println("Result: " + xml2.getStringVal());
    xml2.close(); // free the XMLType result tracked by the Server
}
rset.close();
```

Example 22.a: Using the free() method to free the resources in JDBC

```
SQLXML xml2;
while (rset.next())
{
    xml2 = rset.getSQLXML(1);
    System.out.println("Result: " + xml2.getString());
    xml2.free(); // free the XMLType result tracked by the Server
}
rset.close();
```

When an XMLType is fetched or created using the C API for XML, please make sure to call the corresponding free method to free resources allocated by the server.

- Call OCIXmlDbFreeXmlCtx() to free xmlctx.
- Any xmldocnode obtained via OCISmtExecute, XmlLoadDom, OCIObjectCopy etc. should be freed using XmlFreeDocument().
- If there are OCI Objects or arrays with xml attributes, the top-level object should be freed using OCIObjectFree.
- If an OCI Object is created using XMLType tdo, please free it using OCIObjectFree.

When an XMLType is fetched or created using PL/SQL API for XMLType, remember to call the corresponding free procedure for the document or node instance, when you are through with the instance. Procedure freeDocument frees the document and all of its nodes.

When a variable of XMLType is used in the PL/SQL program, please explicitly set it to NULL to free it:

```
xml_variable := NULL;
```

When a collection of XMLType is used in the PL/SQL program, please use DELETE method to delete the collection elements:

```
collection_xml_variable.DELETE;
```

XQuery Guideline 13: Avoid calling getObject multiple times for xmltype in client program

In JDBC program, please avoid calling getObject() multiple times. Because XMLType object is ref counted, every call to getObject() will increase ref count by one. The call to close() method of XMLType will free the object when the ref count is 1.

Example 23: Avoid calling getObject() twice

Incorrect:

```
Object res = rset.getObject(j);  
  
if( res instanceof XMLType)  
{  
    xml = (XMLType)rset.getObject(j);  
}
```

Correct:

```
Object res = rset.getObject(j);  
  
if( res instanceof XMLType)  
{  
    xml = (XMLType)res;  
}
```

XQuery Guideline 14: Set parameter OPTIMIZER_FEATURE_ENABLE to 11.1.0.6 or above for XQuery optimizations

Storage dependent performance tuning

Recall that Oracle XML DB performs logic rewrite optimization followed by physical rewrite optimization based on XML storage and index by evaluating the XPath expression against the XML document without ever constructing the XML document in memory. This optimization is called XPath rewrite optimization. It is a proper subset of XML query optimization, which also involves optimization of XQuery expressions, such as FLWOR expressions, that are not XPath expressions. XPath rewrite includes **XMLIndex** optimizations, streaming evaluation of binary XML, and rewrite to underlying object-relational or relational structures in the case of structured storage or **XMLType** views over relational data.

XPath rewrite can occur in these contexts (or combinations thereof):

- When **XMLType** data is stored in an object-relational column or table (structured storage) or when an **XMLType** view is built on relational data.
- When you use an **XMLIndex** index.
- When **XMLType** data is stored as binary XML – using streaming evaluation.

All of these items are discussed in the following subsections.

Binary XML

Binary XML storage is used primarily for unstructured data. For this storage, the standard database indexes (B-tree, bitmap) are generally not helpful for accessing particular parts of an XML document. XMLIndex provides a general, XML-specific index that indexes the internal structure of XML data. One of its main purposes is to overcome the indexing limitation presented by binary XML storage. Sometimes when a query cannot use any index, it can still be optimized using the streaming XPath evaluation. This section provides guidance on which indexes to create, and how to write your query to use the streaming XPath evaluation.

Binary XML Streaming Evaluation

The streaming mode of XPath evaluation is used to efficiently evaluate the most common types of XPaths over documents that are stored in Binary XML. This is done by first rewriting the query to collect related XPaths together so that they can be evaluated in a single pass over the document. This type of rewrite is reflected in the output of ‘explain plan’ as ‘XPath EVALUATION’. For example, in the plan for the following query, the XPaths from the SELECT list and the WHERE clause are gathered and evaluated together as columns of the ‘XPath EVALUATION’ step; this is reflected in the predicate information section, which refers to the column corresponding to /PurchaseOrder/Reference.

Example 24: XPath Evaluation in Query Plan

```
SELECT XMLCAST(XMLQuery('$p/PurchaseOrder/@poDate'
                        PASSING OBJECT_VALUE AS "p" RETURNING CONTENT) as DATE)
FROM purchaseorder
WHERE XMLEExists('$p/PurchaseOrder[Reference="123456"]'
                 PASSING OBJECT_VALUE AS "p");
```

Id	Operation	Name	Rows
0	SELECT STATEMENT		1
1	NESTED LOOPS		1
2	TABLE ACCESS FULL	PURCHASEORDER	1
* 3	XPATH EVALUATION		

Predicate Information (identified by operation id):

```
3 - filter("P"."C_01$"='123456')
```

In the following query, all the columns of the XMLTable are evaluated together as part of the ‘XPath EVALUATION’ step:

Example 25: Streaming XPath Evaluation of XMLTable query

```

SELECT li.description, li.lineitem
FROM
  purchaseorder T,
  XMLTable('$p/PurchaseOrder/LineItems/LineItem'
    PASSING OBJECT_VALUE AS "p"
    COLUMNS lineitem    NUMBER          PATH '@ItemNumber',
             description  VARCHAR2(30)   PATH 'Description',
             partid       NUMBER         PATH 'Part/@Id',
             unitprice    NUMBER         PATH 'Part/@UnitPrice',
             quantity     NUMBER         PATH 'Part/@Quantity') li
WHERE li.unitprice > 30 and li.quantity < 20;

```

Id	Operation	Name	Rows
0	SELECT STATEMENT		1
1	NESTED LOOPS		1
2	TABLE ACCESS FULL	PURCHASEORDER	1
* 3	XPATH EVALUATION		

Predicate Information (identified by operation id):

```

3 - filter(CAST("P"."C_01$" AS NUMBER)>30 AND
           CAST("P"."C_02$" AS NUMBER)<20)

```

In general, XPaths involving the child and descendant axes can be evaluated in this mode, but not ones involving reverse axes (like the ancestor axis). Most position-based predicates in XPaths are evaluated in streaming mode in 11gR2. In releases prior to 11gR2, XPaths involving position predicates cannot be evaluated in this mode. In all releases, XPaths with predicates involving `last()`, as well as those with position-based and non-position-based predicates in the same step should be avoided, as these XPaths are not evaluated in streaming mode.

Here are some guidelines on how to write queries to get the best results from streaming XPath evaluation:

Streaming Evaluation Guideline 1: Convert reverse XPath axes to forward axes when possible

In many cases, it is easy to convert an XPath that uses reverse axes to an equivalent one that does not (i.e., uses forward axes only). For example, the following query uses the parent axis (the ‘..’ step) to select nodes that have a child that is named ‘a’ and has an attribute id whose value is ‘abc1’. It can be rewritten to an equivalent query that does not use the parent axis, by including ‘a’ in the predicate (rather than as a separate path step), as shown below.

Example 26: Conversion of reverse axes to forward axes

```
-- Query with reverse axis (cannot be evaluated in streaming mode)
```

```

SELECT XMLQuery('$p/PurchaseOrder/*[a[@id="abc1"]]/..'
                PASSING OBJECT_VALUE AS "p" RETURNING CONTENT)
FROM purchaseorder
WHERE XMLeExists('$p/PurchaseOrder[Reference="123456"]'
                PASSING OBJECT_VALUE AS "p");

-- Equivalent query with no reverse axes (can be evaluated in
streaming mode)
SELECT XMLQuery('$p/PurchaseOrder/*[a[@id="abc1"]]'
                PASSING OBJECT_VALUE AS "p" RETURNING CONTENT)
FROM purchaseorder
WHERE XMLeExists('$p/PurchaseOrder[Reference="123456"]'
                PASSING OBJECT_VALUE AS "p");

```

Streaming Evaluation Guideline 2: For large documents, avoid descendant axis & wild cards if exact (named) path steps can be used

Although XPath with descendant axis & wild cards can be evaluated in streaming mode, they are not as efficient as using just the child axis and named path steps. For example, to get all the line items in a particular purchase order that have a quantity greater than 5, use `/PurchaseOrder/LineItem` instead of `//LineItem`, as shown below.

Example 27: Avoiding descendant axis

```

-- Query with descendant axis
SELECT XMLQuery('$p//LineItem[@quantity > 5]'
                PASSING OBJECT_VALUE AS "p" RETURNING CONTENT)
FROM purchaseorder
WHERE XMLeExists('$p/PurchaseOrder[Reference="123456"]'
                PASSING OBJECT_VALUE AS "p")

-- Query with named path steps (avoiding descendant axis)
SELECT XMLQuery('$p/PurchaseOrder/LineItem[@quantity > 5]'
                PASSING OBJECT_VALUE AS "p" RETURNING CONTENT)
FROM purchaseorder
WHERE XMLeExists('$p/PurchaseOrder[Reference="123456"]'
                PASSING OBJECT_VALUE AS "p")

```

Streaming Evaluation Guideline 3: For DML-heavy workloads, enable caching for writes on the underlying lob column

Note that binary xml tables use a hidden blob column named 'xmldata' to store the encoded xml documents. For workloads that involve a significant amount of DML, enabling caching for writes on this lob column will speed up subsequent queries that use streaming evaluation on the affected documents. The following sql statement enables caching for writes on the purchaseorder table's blob column:

Example 28: Enabling caching for DMLs on binary xml tables

```
ALTER TABLE purchaseorder modify lob (xmldata) (cache);
```

Indexing binary XML

As mentioned above, the indexing solutions in the relational world are not suitable for indexing XML, hence we have a different set of indexes for XML usecases. The different indexes supported in 12gr1 are given below:

- XMLIndex structured component, or Table-based index. This is called the “Structured XMLIndex” for short.
- XMLIndex unstructured component, or Path-based index. . This is called the “Unstructured XMLIndex” for short.
- XML Full Text Index

If you are dealing with large volumes of XML data, you may want to consider taking advantage of the parallelism and partitioning features offered by Oracle. When the base XML is partitioned by range or list partitioning methods, then a corresponding XMLIndex can be created on the XML using the keyword LOCAL. When this is done, the XMLIndex is equi-partitioned with the base table – each partition of the XMLIndex has a 1-1 correspondence with a partition of the base XML. Note that XMLIndex partitioning is only supported on tables that are range or list partitioned.

You can use a PARALLEL clause (with optional degree) when creating or altering an XMLIndex index to ensure that index creation and maintenance are carried out in parallel. If the base table is partitioned or enabled for parallelism, then this can improve the performance for both DML operations (INSERT, UPDATE, DELETE) and index DDL operations (CREATE, ALTER, REBUILD). The degree-of-parallelism (DOP) value specified at the XMLIndex level is also set on each internal table of the XMLIndex – for both structured and unstructured components.

The predicates of path expression, WHERE clause of FLWOR expression, WHERE clause of SQL/XML statement having XMLExists() or XMLTable construct are subject shown in examples below. Such predicate evaluation can be greatly speeded up by using the right XMLIndex. XMLIndex can be used to do both inter-document search (filtering XML document rows stored in the table) and intra-document search (filtering XML document fragment for XML document stored in each row of the table).

Example 29: Examples of where XMLIndex could be used

```

/* XMLExists() with predicate in path expression in SQL WHERE clause:
 * Index can be used to filter rows from table purchaseorder */
SELECT XMLQuery('$po/PurchaseOrder/Requestor'
    PASSING OBJECT_VALUE AS "po" RETURNING CONTENT)
FROM purchaseorder
WHERE
XMLExists('$poPurchaseOrder/LineItems/LineItem/Part[@Quantity = 1]'
    PASSING OBJECT_VALUE AS "po");

/* WHERE clause of XQuery expression.
 * Index can be used to filter rows from table purchaseorder */
SELECT *
FROM XMLTABLE (

```

```

    'for $po in ora:view("purchaseorder")/PurchaseOrder
      where $po/LineItems/LineItem/@ItemNumber="1"
      return $po/Requestor);

/* XMLTable column in SQL WHERE clause. */
SELECT li.description, li.lineitem
FROM purchaseorder,
     XMLTable('/PurchaseOrder/LineItems/LineItem'
              PASSING OBJECT_VALUE
              COLUMNS lineitem    NUMBER          PATH '@ItemNumber',
                       description  VARCHAR2(30)    PATH 'Description',
                       partid       NUMBER          PATH 'Part/@Id',
                       unitprice    NUMBER          PATH 'Part/@UnitPrice',
                       quantity     NUMBER          PATH 'Part/@Quantity') li
WHERE Lineitem = 4567;

/* Predicate in path expression.
 * Index can be used to identify document fragment.
 * Although index cannot be used to identify rows from the purchaseorder
 * because all rows are returned due to no SQL WHERE clause,
 * index can be used to identify Description fragment
 * that satisfies the path predicate from each row of purchasorder.
 * This query is analogous to scalar subquery usage in a select list
 * of a SQL statement where the scalar subquery has its own
 * where clause that can leverage index */

SELECT XMLCAST(XMLQUERY (
  '/PurchaseOrder/LineItems/LineItem[@ItemNumber=1]/Description'
  PASSING object_value RETURNING CONTENT) AS VARCHAR2(4000))
FROM purchaseorder p;

```

Here are some guidelines on which indexes to choose for your usecase.

Index choosing Guideline 1: Use the Structured XMLIndex when XPath's are static, and to answer predicates

If you know your XPath's in advance, the Structured Component of XMLIndex is ideal for your usecase. This will help you get relational performance on your Xqueries for XPath's that have the Structured XMLIndex on them.

You can get optimal performance by using the structured XMLIndex to index the XPath's that appear in the predicates. These predicates can be in the SQL statement, or the predicate of the where clause, or in the Xquery itself, as shown in “Example 29: Examples of where XMLIndex could be used”.

Creating the structured index as depicted in “Example 32: Creating the XMLIndex with Structured Component” can optimize all the queries above.

Index choosing Guideline 2: Use Unstructured XMLIndex when XPath's are not known in advance

Unstructured XMLIndex is useful to index popular subtrees of documents or whole documents, if needed. When a table or column stores a mixed bag of semi-structured documents, or users are allowed to search under whole documents or whole subtrees, unstructured XMLIndex should be used. Note that even when using unstructured XMLIndex, one should try to limit the paths for which the index can be picked using path-subsetting. A smaller index indexing only popular paths will require less storage space and improve query performance as the index will have less disk blocks to scan. XPath paths that are not indexed can still be used in queries and will be processed using Binary XML Streaming Evaluation.

Example 30: Using Unstructured XMLIndex when XPath paths are not known in advance

```
CREATE INDEX po_xmlindex_ix ON purchaseorder (OBJECT_VALUE) INDEXTYPE IS
XDB.XMLINDEX
  PARAMETERS ('PATHS (INCLUDE
              (/PurchaseOrder/LineItems/ /PurchaseOrder/Reference))');

CREATE INDEX msoffice_xix ON MSOfficeDocuments (OBJECT_VALUE) INDEXTYPE
IS XDB.XMLINDEX
  PARAMETERS ('PATHS (INCLUDE (//w:sdt) NAMESPACE MAPPING
              (xmlns:w="http://schemas.openxmlformats.org/wordprocessingml/2006/
              main" ))');
```

Index choosing Guideline 3: Use Full Text Index for full text search requirements

If your application has requirements for full text searching, consider creating a full text index on the base XMLType column.

Example 31: Using Full Text Index to perform full-text search

```
create table po of xmltype;
BEGIN
  CTX_DDL.create_section_group('mysecgroup', 'PATH_SECTION_GROUP');
  CTX_DDL.set_sec_grp_attr('mysecgroup', 'XML_ENABLE', 'T');
END;
/
create index po_ctx_ix on po (object_value) indextype is
ctxsys.context parameters('section group mysecgroup');

select distinct
XMLCast(XMLQuery('$p/PurchaseOrder/ShippingInstructions/address'
  passing po.object_value as "p" returning content)
  as varchar2(256)) "Address"
from po po
where xmlexists('/PurchaseOrder/ShippingInstructions/address[. contains
text "Fortieth" using stemming]' passing po.object_value);
```

Index choosing Guideline 4: Fragment extraction

In the presence of queries that project out XML fragments, the indexing approach depends on the average size of documents:

- If the dataset consists of small to medium size documents, you should use one of the following:

- Either, use the XQuery extension expression (`#ora:xq_proc #`) to indicate XQuery shall be functionally evaluated. *Note: ora:xq_proc gives you fine-grained control -- you can make fragment extraction use xq_proc and predicates use XMLIndex, as long as the predicate XPath is not excluded from the XMLIndex.*
- Or, use Binary XML streaming evaluation – you would need to exclude the path of the fragment to be extracted from your path subsetting Unstructured XMLIndex, if present, to ensure this. *Note: XMLIndex, if present, can still be used to evaluate the predicates, as long as the predicate XPath is not excluded from the XMLIndex.*
- If the dataset contains mostly large documents, use Unstructured XMLIndex to perform fragment extraction by ensuring that the path of the fragment to be extracted is included in the list of paths that are indexed.

Index choosing Guideline 5: Combine different indexes as needed

You can use a combination of the different indexes. For example, if you have a table of technical documents, you can create an XMLIndex with structured component for the title, author and date fields, and with unstructured component for the variable part of the document, and create an Oracle Full Text index to answer text-search queries.

Once you have chosen the right indexes for your use case, please refer to the corresponding section for guidelines on how to get the best performance out of these indexes.

XMLIndex Structured Component

Even though the data in the Binary XML columns may be unstructured, it sometimes contains islands of predictable, structured data. An example is a technical document, with the title, author and date fields. You create and use the structured component of an XMLIndex index for queries that project fixed, structured islands of XML content, even if the surrounding data is relatively unstructured. A structured XMLIndex component organizes such islands in a relational format. It is similar to SQL/XML function XMLTable, and the syntax you use to define the structured component reflects this similarity. The relational tables used to store the indexing data are data-type aware, and each column can be of a different scalar data type. You can thus think of the act of creating the structured component of an XMLIndex index as decomposing a structured portion of your XML data into relational format.

The structured component is a targeted index, and therefore requires careful specification of the XPathS that are to be indexed, along with their data types. But, the benefits of using such an index are significant for queries with statically known XPathS.

Some of the advantages of using the structured component are listed below:

1. *Type-Aware, Relational-Style Searches* - The path table of the unstructured XMLIndex component stores values of all indexed nodes in the same column, irrespective of the actual schema-type of that value. Number values, date values, and string values are stored in the same column. Therefore, secondary indexes on this single value column do not provide proper statistics information, and therefore are not easily picked up by the relational Cost Based Optimizer (CBO) within Oracle. The structured component of an XMLIndex has the ability to separate values by type and by path into different columns, and therefore can provide very specific relational-style statistics to the relational Cost Based Optimizer, on which the XMLIndex is built.
2. *Support for Composite B-Tree and Bitmap Indexes* – Since values from various XPathS are stored in the same column of the path table of the unstructured component, it is not possible to create bitmap indexes or composite indexes, even though the query or the data suggest that such indexes are appropriate. On the other hand, an internal table of a structured XMLIndex can store values from different XPathS in separate columns, thereby making it possible to create composite B-Tree and bitmap indexes.
3. *No Sub-Query in SQL Predicate* - When the value-search criterion appears in the WHERE clause of a SQL statement, then rewrite using the unstructured XMLIndex component results in a sub-query in the WHERE clause. Such a sub-query cannot be folded into the top-level query, leading to poor execution plans. This problem is particularly an issue when the XML does not come with an XML schema. On the other hand, when

structured XMLIndex is used, a predicate in the WHERE clause becomes column-level checks on the structured XMLIndex tables.

4. *Indexing for BI-Style Queries* – SQL constructs such as order-by, group-by, window, etc., enable powerful business intelligence queries over relational data. Applications using order-by, group-by, window, etc., on values within XML data can get relational performance by using structured XMLIndex, since the queries can be rewritten to order-by, group-by, window, etc., over relational table columns. This is accomplished as follows: XMLTable allows values in XML to be projected out as a virtual table. A query that uses the XMLTable function can be rewritten to simple access of the relational tables of a structured XMLIndex. This means that order-by, group-by, window, etc., operating on columns of the virtual table are translated to order-by, group-by, window, etc., operating on the corresponding physical columns of the structured XMLIndex tables. When unstructured XMLIndex is used, the order-by, group-by, window, etc., functions are evaluated as sub-queries over the path table, and therefore performs orders of magnitude slower than when structured XMLIndex is used.

The example below shows how to create structured XMLIndex. It uses the Purchase Order schema, which has a collection called “LineItem”. For each XML node matching the row pattern /PurchaseOrder/LineItems/LineItem, this XMLIndex projects out in its relational index table 5 columns – the values of these nodes are the values of nodes matching relative XPath’s @ItemNumber, Description, Part/@Id, Part/@UnitPrice, and Part/@Quantity. The internal index table will have as many rows for each XML document as the number of LineItem nodes within the document. The index DDL specifies the name of the table (lineitem_tab in this case), the names of the 5 columns, and the SQL data types of these 5 columns.

Example 32: Creating the XMLIndex with Structured Component

```
CREATE INDEX po_struct ON purchaseorder (OBJECT_VALUE)
INDEXTYPE IS XDB.XMLIndex
PARAMETERS (
'XMLTable lineitem_tab ''/PurchaseOrder/LineItems/LineItem''
  COLUMNS lineitem      NUMBER          PATH '@ItemNumber'',
           description  VARCHAR2(30)    PATH 'Description'',
           partid       NUMBER          PATH 'Part/@Id'',
           unitprice    NUMBER          PATH 'Part/@UnitPrice'',
           quantity     NUMBER          PATH 'Part/@Quantity''');
```

Below are the guidelines on how to get the best performance out of your structured XMLIndex.

Structured Index Guideline 1: Use Structured Index instead of multiple functional indexes and/or virtual columns

In XML usecases where user wants to project out several relational key columns of XML so that they can build B-tree indexes over these columns for quick search, structured XMLIndex is ideal.

Structured XMLIndex projects out one relational table capturing all the key relational columns for efficient search, instead of relying on multiple virtual columns (VC) that are inefficient. These structured XMLIndex columns are efficiently populated in a single scan of the input base document - something that cannot be done with virtual columns. Also, the structured XMLIndex based approach works in cases where the XML has collections, whereas the VC based approach cannot be used when the projected value is within an XML collection.

Structured Index Guideline 2: Make Index and Query datatypes correspond

The relational tables that are used for an XMLIndex structured component use SQL data types. XQuery expressions that are used in queries use XML data types (XML Schema data types and XQuery data types). XQuery typing rules can automatically change the data type of a subexpression, to ensure coherence and type-checking. For example, if a document that is queried using XPath expression `/PurchaseOrder/LineItem[@ItemNumber = 25]` is not XML schema-based, then the subexpression `@ItemNumber` is `xs:untypedAtomic`, and it is then automatically cast to `xs:double` by the XQuery `=` comparison operator. To index this data using an XMLIndex structured component you must use `BINARY_DOUBLE` as the SQL data type.

This is a general rule. For an XMLIndex index with structured component to apply to a query, the data types must correspond. Table 2 in Guideline 7: “Using proper XQuery and SQL Typing”

If the XML and SQL data types involved do not have a built-in one-to-one correspondence, then you must make them correspond (according to Table 2), in order for the index to be picked up for your query. There are two ways you can do this:

- Make the index correspond to the query – Define (or redefine) the column in the structured index component, so that it corresponds to the XML data type. For example, if a query that you want to index uses the XML data type `xs:double`, then define the index to use the corresponding SQL data type, `BINARY_DOUBLE`.
- Make the query correspond to the index – In your query, explicitly cast the relevant parts of an XQuery expression to data types that correspond to the SQL data types used in the index content table.

Structured Index Guideline 3: Use XMLTable views with corresponding index, e.g BI style queries

Since the structured component of XMLIndex is built on the idea of an XMLTable, such an index fits nicely for usecases where this relational paradigm is applicable. For application developers who want a relational access paradigm, one or more relational views built on XMLTable should be created. The XMLTable function provides a way to expose key values from within XML as relational columns. Querying of XML in many usecases can be hidden within the definitions of relational views that use the XMLTable function, making it easier for

XML to penetrate into the world of application developers who are familiar with SQL and want to be spared the complexity of XPath/XQuery. In such cases, the structured XMLIndex definition will match the definitions of the relational views.

The example below shows how XMLTable() provides a relational table abstraction over XML, and the next example shows how to create a corresponding view, and example 50 shows the corresponding index for it.

Example 33: XMLTable Provides a Virtual Table Abstraction over XML

```
SELECT lines.lineitem ,
       lines.description,
       lines.partid ,
       lines.unitprice ,
       lines.quantity
FROM   purchaseorder,
       XMLTable ('/PurchaseOrder/LineItems/LineItem'
PASSING OBJECT_VALUE
COLUMNS lineitem      NUMBER      PATH '@ItemNumber',
description VARCHAR2(30) PATH 'Description',
partid      NUMBER      PATH 'Part/@Id',
unitprice   NUMBER      PATH 'Part/@UnitPrice',
quantity    NUMBER      PATH 'Part/@Quantity') lines;
```

```
LINEITEM DESCRIPTION PARTID UNITPRICE QUANTITY
-----
11 Orphic Trilogy 37429148327 80 3
22 Dreyer Box Set 37429158425 80 4
11 Dreyer Box Set 37429158425 80 3
```

Example 34: Relational View Using XMLTable, and corresponding structured XMLIndex

```
CREATE VIEW lineitems_v
(lineitem, description, partid, unitprice, quantity)
AS SELECT
       lines.lineitem, lines.description, lines.partid,
       lines.unitprice, lines.quantity
FROM   purchaseorder,
       XMLTable ('/PurchaseOrder/LineItems/LineItem'
PASSING OBJECT_VALUE
COLUMNS lineitem      NUMBER      PATH '@ItemNumber',
description VARCHAR2(30) PATH 'Description',
partid      NUMBER      PATH 'Part/@Id',
unitprice   NUMBER      PATH 'Part/@UnitPrice',
quantity    NUMBER      PATH 'Part/@Quantity'
) lines;
```

One common usecase for this is that of BI-style queries. SQL constructs such as order-by, group-by, window, etc., enable powerful business intelligence queries over relational data. XMLTable allows values in XML to be projected out as a virtual table. Order-by, group-by, window, etc., can operate on columns of the virtual table. Structured XMLIndex internally organizes its storage tables in a manner that reflects the virtual table(s) exposed by XMLTable. Therefore, structured XMLIndex is well suited for indexing XML data in a way that makes such

XMLTable based queries very efficient. A query that uses the XMLTable function can be rewritten to simple access of the relational tables of a structured XMLIndex. This means that order-by, group-by, window, etc., operating on columns of the virtual table are translated to order-by, group-by, window, etc., operating on the corresponding physical columns of the structured XMLIndex tables.

We recommend that the user create relational views over XML using XMLTable, where the views project all columns of interest to the BI application. Application queries should be written against these relational views. If structured XMLIndex is created in 1-1 correspondence to these views, Oracle RDBMS will make sure that queries over the views are seamlessly translated into queries over the relational tables of the structured XMLIndex, thereby giving relational performance.

Structured Index Guideline 4: Create Secondary Indexes, especially for predicates

“Example 32: Creating the XMLIndex with Structured Component” creates relational table `lineitem_tab` under the covers. To get good performance for value-based searches, it is important that the user create secondary indexes on the index table. This is illustrated in the example below.

Example 35: Creating Secondary Indexes on Structured XMLIndex Tables

```
CREATE INDEX li_itemnum_idx      ON lineitem_tab(lineitem);
CREATE INDEX li_desc_idx       ON lineitem_tab(description);
CREATE INDEX li_partid_idx     ON lineitem_tab(partid);
CREATE INDEX li_uprice_idx     ON lineitem_tab(unitprice);
CREATE INDEX li_quantity_idx   ON lineitem_tab(quantity);
```

Composite B-Tree indexes, bitmap indexes and domain indexes (e.g., Oracle Text) can also be created on the index table.

Example 36: Creating Oracle Text Index on Structured XMLIndex Table

```
CREATE INDEX li_desc_ctx_idx ON lineitem_tab(description)
indextype is ctxsys.context;
```

It is the responsibility of the user to create these secondary indexes. No secondary index is created automatically by the system for the structured XMLIndex component, as the user is the best judge of what secondary index best suites his needs. Once the secondary indexes are created, the user should gather statistics on the base table so that the optimizer can pick up the indexes.

If a query uses a particular XPath in a predicate, including the SQL WHERE clause, then creating a secondary index on the corresponding column of the structured XMLIndex table is highly recommended.

Structured Index Guideline 5: Check the execution plan to see if structured index is used

After creating the necessary indexes to speed up your queries, you need to verify that the execution plan is indeed picking up the index. For example, let's say you have created the structured XMLIndex as depicted in “Example 32: Creating the XMLIndex with Structured Component”, and secondary indexes as depicted in Example 102. Then you should run an explain plan on your query, as illustrated in the example below:

Example 37: Using Explain Plan to determine that the index is picked up

```
EXPLAIN PLAN FOR
SELECT XMLCAST(XMLQUERY( '/PurchaseOrder/Requestor' PASSING object_value
RETURNING CONTENT) AS VARCHAR2(4000))
FROM purchaseorder p
WHERE
XMLExists('/PurchaseOrder/LineItems/LineItem[xs:decimal(@ItemNumber)=1]'
PASSING object_value);
```

Explained.

```
SQL> select Id, Operation, Name from table(dbms_xplan.display);
```

```
PLAN_TABLE_OUTPUT
-----
Plan hash value: 2801523227
-----
| Id | Operation                                | Name                |
-----|-----|-----|
| 0  | SELECT STATEMENT                          |                      |
| 1  | NESTED LOOPS SEMI                          |                      |
| 2  | TABLE ACCESS FULL                          | PURCHASEORDER       |
|* 3  | TABLE ACCESS BY INDEX ROWID                | LINEITEM_TAB        |
|* 4  | INDEX RANGE SCAN                            | LI_ITEMNUM_IDX      |
-----|-----|-----|
```

The execution plan shows that the query gets rewritten to use the structured index storage table LINEITEM_TAB and the secondary index LI_ITEMNUM_IDX.

Structured Index Guideline 6: Indexing Master-Detail relationships

In cases where the structured islands have a master-detail kind of relationship, structured XMLIndex provides a way to capture each structured island as a relational table, with a primary-foreign key relationship between the tables. Here are definitions of such a master-detail view, and its corresponding structured XMLIndex:

Example 38: Relational View with Master-Detail Relationship

```
CREATE OR REPLACE VIEW purchaseorder_detail_view AS
SELECT po.reference, li.*
FROM purchaseorder p,
XMLTable('/PurchaseOrder' PASSING p.OBJECT_VALUE
COLUMNS
reference VARCHAR2(30) PATH 'Reference',
lineitem XMLType PATH 'LineItems/LineItem') po,
XMLTable('/LineItem' PASSING po.lineitem
```

```

COLUMNS
    itemno          NUMBER(38)    PATH '@ItemNumber',
    description    VARCHAR2(256) PATH 'Description',
    partno         VARCHAR2(14)   PATH 'Part/@Id',
    quantity       NUMBER(12, 2)  PATH 'Part/@Quantity',
    unitprice      NUMBER(8, 4)   PATH 'Part/@UnitPrice') li;

```

Example 39: Structured XMLIndex to Index Master-Detail Relationship

```

CREATE INDEX po_struct ON po_tab (OBJECT_VALUE)
INDEXTYPE IS XDB.XMLIndex
PARAMETERS ('XMLTable po_ptab
    XMLNAMESPACES (DEFAULT 'http://www.example.com/po'),
    '/purchaseOrder'
    COLUMNS orderdate DATE          PATH '@orderDate',
             Id          BINARY_DOUBLE PATH '@id',
             items       XMLType      PATH 'items/item'
VIRTUAL
    XMLTable li_tab
    XMLNAMESPACES (DEFAULT 'http://www.example.com/po'),
    '/item' PASSING items
    COLUMNS partnum   VARCHAR2(15)  PATH '@partNum',
             description CLOB         PATH 'productName',
             usprice   BINARY_DOUBLE PATH 'USPrice',
             shipdat   DATE          PATH 'shipDate');

```

Structured Index Guideline 7: Split fragment extraction and value search between SELECT and WHERE clause

Instead of using a single XQuery for fragment extraction as well as for value search, use XMLQuery() in the SELECT clause for fragment extraction and use XMLExists() in the WHERE clause for value search. By doing this separation, we are able to make structured xmlindex be picked up for value search, while binary XML streaming is used for fragment extraction.

Example 40: Splitting fragment extraction and value search

In this example, Query 1 is a better formulation than Query 2 when following XMLIndex is present:

Index definition:

```

CREATE TABLE XML_TEST (XML_DOC XMLType)
XMLType XML_DOC STORE AS BINARY XML;

CREATE INDEX XML_TEST_IX ON XML_TEST (XML_DOC)
INDEXTYPE IS XDB.XMLIndex
PARAMETERS ('GROUP XML_TEST_G XMLTable XML_TEST_X
    XMLNAMESPACES ('http://example.com/metadata' as "m"),
    '/m:object' COLUMNS
    TENANT VARCHAR(100) PATH 'm:meta/m:tenant',
    ID VARCHAR(250) PATH 'm:meta/m:id');

CREATE INDEX XML_TEST_IX_1 ON XML_TEST_X(TENANT, ID);

```

Query 1: Better

```

SELECT
XMLQUERY('declare namespace m="http://example.com/metadata";
         for $obj in $doc/m:object
         return <m:object>
             {$obj/m:meta/m:id}{$obj/m:meta/m:tenant}
         </m:object>')
         passing T.XML_DOC as "doc" returning content)
FROM XML_TEST T
WHERE
XMLEXISTS('declare namespace m="http://example.com/metadata";
          $doc/m:object[m:meta/m:tenant=$tenant
                       and m:meta/m:id=$id]')
         passing T.XML_DOC as "doc",
          'tenant5' as "tenant",
          'id_555' as "id");

```

Query 2: Avoid

```

SELECT X.XML_DOC
FROM XML_TEST T,
XMLTABLE(
XMLNAMESPACES('http://example.com/metadata' as "m"),
'for $obj in /m:object
 where $obj/m:meta/m:tenant="tenant5" and
       $obj/m:meta/m:id="id_5555"
 return <m:object>
       {$obj/m:meta/m:id}{$obj/m:meta/m:tenant}
 </m:object>')
PASSING T.XML_DOC COLUMNS XML_DOC XMLTYPE PATH '.') X;

```

Structured Index Guideline 8: For ordering query results, use SQL ORDER BY along with XMLTable

Instead of using XQuery ORDER BY clause, use XMLTable to project out the key by which to order and then use SQL ORDER BY. In the example below, the query shows fragment extraction together with value search. Fragments are ordered by tenant, id which are projected out in the XMLTABLE() clause:

Example 41: Using SQL order by

```

SELECT XMLQUERY('declare namespace
m="http://example.com/metadata";
         for $obj in $doc/m:object
         return <m:object>
             {$obj/m:meta/m:id} {$obj/m:meta/m:tenant}
         </m:object>')
         passing T.XML_DOC as "doc" returning content)
FROM XML_TEST T,
XMLTABLE(XMLNAMESPACES('http://example.com/metadata'
                        as "m"),

```

```
        '$doc/m:object' PASSING T.XML_DOC as "doc"
        COLUMNS
            tenant VARCHAR(100) PATH 'm:meta/m:tenant',
            id VARCHAR(250) PATH 'm:meta/m:id'
        ) tt
WHERE
    XMLEXISTS('declare namespace
m="http://example.com/metadata";
    $doc/m:object[m:meta/m:tenant=$tenant]'
    passing T.XML_DOC as "doc", 'tenant5' as "tenant")
ORDER BY tt.tenant, tt.id;
```

XMLIndex Unstructured Component

The XMLIndex Structured component is not suited for documents that involve little structure or queries that extract XML fragments. The XMLIndex unstructured component is general and relatively untargeted.

Unlike a B-tree index, which you define for a specific database column that represents an individual XML element or attribute, or the XMLIndex structured component, which applies to specific, structured document parts, the unstructured component of an XMLIndex index is, by default, very general. Unless you specify a more narrow focus by detailing specific XPath expressions to use or not to use in indexing, the XMLIndex unstructured component applies to all possible XPath expressions for your XML data.

The XMLIndex unstructured component uses a path table and a set of (local) secondary indexes on the path table, which implement the logical parts described above. Two secondary indexes are created automatically:

- A pikey index, which implements the logical indexes for both path and order.
- A real value index, which implements the logical value index.

You can modify these two indexes or create additional secondary indexes. The path table and its secondary indexes are all owned by the owner of the base table upon which the XMLIndex index is created. Though you can restrict an unstructured component to apply only to certain XPath subsets, its path table indexes node content can be of different scalar types. This can require you to create multiple secondary indexes on the VALUE column to deal with the different data types.

The PIKEY index handles paths and order relationships together, which gives the best performance in most cases. If you find in some particular case that the value index is not picked up when think it should be, you can replace the pikey index with separate indexes for the paths and order relationships. Such (optional) indexes are called path id and order key indexes, respectively.

The path table contains one row for each indexed node in the XML document. For each indexed node, the path table stores:

- The corresponding rowid of the table that stores the document.
- A locator, which provides fast access to the corresponding document fragment. For binary XML storage of XML schema-based data, it also stores data-type information.
- An order key, to record the hierarchical position of the node in the document. You can think of this as a Dewey decimal key like that used in library cataloging and Internet protocol SNMP. In such a system, the key 3.21.5 represents the node position of the fifth child of the twenty-first child of the third child of the document root node.

- An identifier that represents an XPath path to the node.
- The effective text value of the node.

Below are some guidelines on how to get the best performance out of the XMLIndex Unstructured component.

Unstructured XMLIndex Guideline 1: Check the Execution Plan to see if the XMLIndex Unstructured Component is used

At the time of creating an XMLIndex, one can specify the name of the Path Table to be used in the PARAMETERS clause. If the name of the Path Table is not specified, a system generated name will be given to the path table. Users can figure out the name of the path table by querying on view *user_xml_indexes*.

A plan that picks up unstructured index should use a scan on the Path Table for the index. It should be an INDEX RANGE SCAN with one of the secondary indexes. Names of the secondary indexes are also visible in view *user_xml_indexes*. The part of the plan in the example below shows the use of path table and its secondary index:

Example 42: Plan showing the use of path table and secondary index

```
|*2| TABLE ACCESS BY INDEX ROWID| MY_PATH_TABLE|
|*3| INDEX RANGE SCAN          | SYS67616_PO_XMLINDE_PIKEY_IX |
    2 - filter(SYS_XMLI_LOC_ISNODE("SYS_P0"."LOCATOR")=1)
    3 - access("SYS_P0"."RID"=:B1 AND "SYS_P0"."PATHID"=HEXTORAW('6F7F') )
```

Unstructured XMLIndex Guideline 2: When to drop PIKEY index in favor of ORDERKEY & PATHID index

PIKEY index is the most useful for complex XQueries and least useful for very simple SQL/XML based formulations.

If you find in some particular case that the value index is not picked up when think it should be, you can replace the PIKEY index with separate indexes for the paths and order relationships. Such (optional) indexes are called PATH ID and ORDER KEY indexes, respectively. They can be specified in the parameter clauses “PATH ID INDEX” and “ORDER KEY INDEX” respectively in DDLs.

The example below illustrates that if an application uses most queries of the following form, it should consider switching to PATH ID and ORDER KEY indexes:

Example 43: When to use which secondary index

```
xmlexists( /a/b[c=5] ) => only VALUE index is needed
xmlexists( /a/b/c ) => simple PATH ID index would offer better performance.
```

```
select v.c1, v.c2,... from t,
```

`xmltable(/a/b columns c1 path 'c', c2, ...) v => PATH ID, ORDER KEY` will probably offer better performance than PIKEY since each of the select list items will benefit from an ORDER KEY based join rather than a PATH ID based one.

Unstructured XMLIndex Guideline 3: How to use path-subsetting -- smaller index means faster queries

If the queries access only parts of the document, these parts can be specified using the PATHS parameter in the PARAMETER clause. This will ensure that the Path Table and the secondary indexes are smaller, and will lead to faster Query and DML performance. To ensure that XMLIndex is indeed picked up for the queries, check the execution plan for the parts specified in Guideline 1.

Unstructured XMLIndex Guideline 4: Using path-subsetting to choose streaming vs index execution

For Binary XML storage, streaming evaluation is very well optimized for fragment extraction. Path Subsetting can be used to take advantage of the best of XMLIndex and Streaming Evaluation. Users may use Path Subsetting to INCLUDE only the XPaths that are in the predicates or EXCLUDE any XPaths that are in the results. For example, queries on Purchaseorder table have predicates on various children of LineItem elements and return the BillTo and ShipTo Address elements. In such scenarios, we can define an XMLIndex that includes LineItem elements.

Unstructured XMLIndex Guideline 5: Using NO_XMLINDEX_REWRITE_IN_SELECT hint

In 11.2.0.3, for a large collection of small documents, users can use the hint NO_XMLINDEX_REWRITE_IN_SELECT for selective queries to let the index be used for filtering in the where clause while letting the select list be evaluated via streaming. For example, the following query from the Maturity benchmark will use UXI, if present, for the XMLEXISTS clause but will not use UXI to extract the fragment \$cust/Accounts in the XMLQuery return clause:

```
select /*+ NO_XMLINDEX_REWRITE_IN_SELECT*/ XMLQUERY('declare default
element namespace "http://tpox-benchmark.com/custacc"; for $cust in
$cadoc/Customer

return $cust/Accounts'

PASSING v.object_value AS "cadoc" returning content)

FROM CUSTACC v

WHERE XMLEXISTS
```

```

('declare default element namespace "http://tpox-
benchmark.com/custacc"; $cadoc/Customer[@id < 1060]' PASSING
v.object_value AS "cadoc")

```

Unstructured XMLIndex Guideline 6: Creating datatype aware VALUE indexes by making index and query datatypes correspond

The DBMS_XMLINDEX package provides functions to create typed indexes on the value column of the path table:

1. CreateNumberIndex
2. CreateDateIndex

Based on the query requirements, if the predicates are on number or date values, the respective secondary index should be created. Please refer to XQuery Guideline 7.

For typical non-schema based XML document search using numeric value comparison, such as '/company[id = 3456]', XQuery semantics entails this as xs:double() type comparison. Therefore, createNumberIndex() is invoked passing 'double' as *xmlypename* parameter as follows:

```

DBMS_XMLINDEX.CREATENUMBERINDEX('SCOTT', 'PO_XMLINDEX_IX',
                                'PO_DOUBLE_NUM_IX', '', 'double');

```

The following table shows the value of *xmlypename* parameter to be used during secondary index creation.

TABLE 4. XML AND SQL DATA TYPE CORRESPONDENCE FOR XMLINDEX

XML DATA TYPE	SQL DATA TYPE	XMLTYPE NAME FOR CREATE NUMBER INDEX	XMLTYPE NAME FOR CREATE DATE INDEX
xs:integer, xs:decimal	INTEGER or NUMBER	integer, decimal	
xs:double	BINARY_DOUBLE	double	
xs:float	BINARY_FLOAT	float	
xs:date	DATE, TIMESTAMP WITH TIMEZONE		date
xs:time, xs:dateTime	TIMESTAMP, TIMESTAMP WITH TIMEZONE		time, datetime
xs:dayTimeDuration	INTERVAL DAY TO		Not available

xs:yearMonthDuration	SECOND INTERVAL YEAR TO MONTH		Not available
----------------------	-------------------------------------	--	---------------

Unstructured XMLIndex Guideline 7: XPath Expressions not indexed by Path Subsetted XMLIndex

Unstructured XMLIndex indexes everything in the XML document but a Path Subsetted index indexes only specific subtrees of a document. Thus, a Path Subsetted index cannot be used for queries with parent axis,

Unstructured XMLIndex Guideline 8: Be specific in the XPath (avoid //, /*)

Even though Unstructured XMLIndex indexes all nodes and hence can be used to query all possible XPaths, query plans with XPaths that have // and /* in the middle need a join of the Path Table with itself or with the Binary XML token tables causing sub-optimal plans. If the structure of the document is known, XPaths should be more specific and if the structure of the documents is not known, avoid prefixing // with ancestor elements. E.g. use //c and not /a/b//c, provided these return the same result set.

Note that each descendant axis access involves a lookup of the suffix path in the token table. This involves the use of SYS_PATH_REVERSE operator on the PATH column in the token table. A typical plan will have the following constructs:

```

14|      TABLE ACCESS BY INDEX ROWID| X$PT48E463W8DU8V6E0G741BMDT9 |
15|      INDEX RANGE SCAN              | X$PR48E463W8DU8V6E0G741BMDT9 |
16|      INDEX RANGE SCAN              | SYS67616_POIX_PIKEY_IX       |
17|      TABLE ACCESS BY INDEX ROWID | SYS67616_POIX_PATH_TABLE     |
18|      TABLE ACCESS FULL            | PO                             |

```

Predicate Information (identified by operation id):

```

5 - access(SYS_PATH_REVERSE("PATH")>=HEXTORAW('021D34') AND
      SYS_PATH_REVERSE("PATH")<HEXTORAW('021D34FF') )
6 - access("SYS_P0"."RID"=:B1 AND "SYS_P0"."PATHID"="ID")
7 - filter(SYS_XMLI_LOC_ISNODE("SYS_P0"."LOCATOR")=1)

```

Unstructured XMLIndex Guideline 9: Reduce the number of expressions in the from clause (avoid Path Table join with itself)

The virtual tables created using the XMLTable clause in the FROM list should be reduced in number, if possible. Each such XMLTable leads to a Path Table in the XMLIndex rewritten query. The example below shows which kinds of queries perform better:

Example 44: Queries to use and avoid

Use this form of query:

```

SELECT li.description
FROM po_clob p,
     XMLTable('PurchaseOrder/LineItems/LineItem' PASSING p.OBJECT_VALUE
              COLUMNS description VARCHAR2(256) PATH 'Description') li;

```

Avoid this form of query, if possible:

```

SELECT li.description
FROM po_clob p,
     XMLTable('PurchaseOrder/LineItems' PASSING p.OBJECT_VALUE) ls,
     XMLTable('LineItems/LineItem' PASSING ls.OBJECT_VALUE
              COLUMNS description VARCHAR2(256) PATH 'Description') li;

```

Unstructured XMLIndex Guideline 10: Use of an index on sys_orderkey_depth

If you use an XPath expression in a query to drill down inside a virtual table (created, for example, using SQL/XML function XMLTable), then create a secondary index on the order key of the path table using Oracle SQL function *sys_orderkey_depth*. The example below shows such a query; the selection navigates to element Description inside virtual line-item table li.

Such queries are evaluated using function *sys_orderkey_depth*, which returns the depth of the order-key value. Because the order index uses two columns, the index needed is a composite index over columns ORDER_KEY and RID, as well as over function *sys_orderkey_depth* applied to the ORDER_KEY value.

Example 45: Query drilling down inside virtual table, and index to create for it

```

SELECT li.description
FROM po_clob p,
     XMLTable('PurchaseOrder/LineItems/LineItem' PASSING p.OBJECT_VALUE
              COLUMNS description VARCHAR2(256) PATH 'Description') li;
CREATE INDEX depth_ix ON my_path_table
  (RID, sys_orderkey_depth(ORDER_KEY), ORDER_KEY);

```

Unstructured XMLIndex Guideline 11: Old snapshot queries might be slow

In case the XMLIndex is created with an ASYNC option with STALE=TRUE and the index has not been synced for a long time, the query performance degrades as the query has to access the data from the last-synced snapshot. If your system suddenly shows slow performance with the same query plans, this might be the reason.

Unstructured XMLIndex Guideline 12: Avoid the usage of text() in path expression

Unstructured XMLIndex rewrite generates sub-optimal query plans for path expression containing text(), especially inside a predicate. It is recommended to specify element name test or context item expression instead. E.g. use /a/b[c="foo"] not /a/b[c/text()="foo"], use /a/b[.="bar"] not /a/b[text()="bar"].

Full Text Index

Besides accessing XML nodes such as elements and attributes, it is sometimes important to provide fast access to particular passages of text within XML text nodes. This is the purpose of Oracle Text indexes: they index full-text strings. Full-text indexing is particularly useful for document-centric applications, which often contain a mix of XML elements and text-node content. Full-text searching can often be made more powerful and more focused, by combining it with structural XML searching, that is, by restricting it to certain parts of an XML document, which are identified by using XPath expressions.

An XML-enabled Oracle Text CONTEXT index, also referred to as an XML full-text index, created on an XMLType column enables XQuery Full Text (XQFT) queries on XMLType data..

To create an XML-enabled Oracle Text index you must be granted database role CTXAPP. More generally, this role is needed to create Oracle Text indexes, to set Oracle Text index preferences, or to use Oracle Text PL/SQL packages.

Before creating the index, you must create an Oracle Text path section group and set its xml_enable attribute to true (t). This makes the path section group XML-aware.

For best performance, create an index preference of type BASIC_STORAGE in the Oracle Text data dictionary, specifying the following attributes:

- **D_TABLE_CLAUSE** – Specify SECUREFILE storage for column DOC of index data table \$D, which contains information about the structure of your XML documents. Specify caching and medium compression.
- **I_TABLE_CLAUSE** – Specify SECUREFILE storage for column TOKEN_INFO of index data table \$I, which contains information about full-text tokens and their occurrences in the indexed documents. Specify caching (but not compression).

The example below shows how to create an XML-enabled Oracle Text index on an XMLType column.

Example 46: Creating an XML-enabled Oracle Text Index

```
BEGIN
  CTX_DDL.create_section_group('mysecgroup', 'PATH_SECTION_GROUP');
  CTX_DDL.set_sec_grp_attr('mysecgroup', 'XML_ENABLE', 'T');
  CTX_DDL.create_preference('mypref', 'BASIC_STORAGE');
  CTX_DDL.set_attribute('mypref',
                       'D_TABLE_CLAUSE',
                       'TABLESPACE my_ts
```

```

        LOB(DOC) STORE AS SECUREFILE
        (TABLESPACE my_ts COMPRESS MEDIUM CACHE)');

CTX_DDL.set_attribute('mypref',
        'I_TABLE_CLAUSE',
        'TABLESPACE my_ts
        LOB(TOKEN_INFO) STORE AS SECUREFILE
        (TABLESPACE my_ts NOCOMPRESS CACHE)');

END;

/
CREATE INDEX po_ctx_idx ON po_binxml(OBJECT_VALUE)

```

```

    INDEXTYPE IS CTXSYS.CONTEXT
    PARAMETERS('storage mypref section group mysecgroup');

```

Index preference BASIC_STORAGE specifies the tablespace and creation parameters for the database tables and indexes that constitute an Oracle Text index.

Example 56 queries the data to retrieve the Description elements whose text contains both Big and Street, in that order.

Example 56: XQuery Full Text Query

```

SELECT XMLQuery('for $i in /PurchaseOrder/LineItems/LineItem/Description
        where $i[. contains text "Big" ftand "Street"]
        return <Title>{$i}</Title>'
        PASSING OBJECT_VALUE RETURNING CONTENT)

FROM po_binxml

WHERE XMLEExists('/PurchaseOrder/LineItems/LineItem/Description
        [. contains text "Big" ftand "Street"]'
        PASSING OBJECT_VALUE);

```

Example 57 shows the execution plan for the query, which indicates that index po_ctx_idx is picked up.

Example 57: Execution Plan for XQuery Full Text Query

```

-----
-----
| Id  | Operation                                | Name          | Rows  | Bytes | Cost (%CPU) |
Time  |

```

```

-----
-----
| 0 | SELECT STATEMENT          |          | 1 | 2014 | 4 (0) |
00:00:01 |
| 1 | TABLE ACCESS BY INDEX ROWID| PO_BINXML | 1 | 2014 | 4 (0) |
00:00:01 |
|* 2 | DOMAIN INDEX                | PO_CTX_IDX |   |   | 4 (0) |
00:00:01 |
-----
-----

```

Predicate Information (identified by operation id):

```

2 - access ("CTXSYS"."CONTAINS" (SYS_MAKEXML(0,"XMLDATA"), '<query><textquery
      grammar="CONTEXT" lang="english"> ( ( {Big} ) and ( {Street} ) )
INPATH(/PurchaseOrder/LineItems/LineItem/Description)</textquery></query>')>0)

```

To get the best performance for your full text queries, follow the guidelines given below:

Full Text Index Guideline 1: What To Do If an XML Full-Text Index Is Not Picked Up

If you use an XQuery full-text predicate in an XMLExists expression within a SQL WHERE clause, but you do not create an XML-enabled Oracle Text index or the index cannot be used for some reason, then compile-time error ORA-18177 is raised.

If this error is raised then your execution plan does not indicate that the index is picked up: in the plan you do not see operation DOMAIN INDEX followed by the name of the index.

In that case, try to change your query to enable the index to be used. The following conditions must both apply for the index to be picked up:

- The expression that computes the XML nodes for the search context must be an XPath expression whose steps are only along forward and descendent axes. Example 58 shows an XMLExists query with following-sibling axis, then compile-time error ORA-18177 is raised.

Example 58: XQuery Full Text Query with following-sibling Axis: Error ORA-18177

```

SELECT *
from po

```

```
where XMLExists('/PurchaseOrder/LineItems/LineItem[Description contains text
"travels" and following-sibling::LineItem contains text "Diary"]' PASSING
OBJECT_VALUE );
```

ERROR at line 3:

```
ORA-18177: XQuery full text expression
'/PurchaseOrder/LineItems/LineItem[Description contains text "travels" and
following-sibling::LineItem contains text "Diary"]' cannot be evaluated using
XML full-text index
```

- You can pass only one XMLType instance as a SQL expression in the PASSING clause of SQL/XML function XMLExists, and each of the other, non-XMLType SQL expressions in that clause must be either a compile-time constant of a SQL built-in data type or a bind variable that is bound to an instance of such a data type. In the following example, the non-XMLType SQL expression is not a compile-time constant, so compile-time error ORA-18177 is raised.

Example 59: XQuery Full Text Query with a dynamic non-XMLType SQL expression: Error ORA-18177

```
select *
from po p
where xmlexists('/purchaseOrder/items/item[productName contains text {$x}]'
    passing p.object_value,
    (select x from just_scal) as "x");
```

ERROR at line 1:

```
ORA-18177: XQuery full text expression '/purchaseOrder/items/item[productName
contains text {$x}]' cannot be evaluated using XML full-text index
```

Example 60 shows the modified query that will enable the index to be used.

Example 60: Revised Query of Example 59 with a constant non-XMLType SQL expression

```
select *
from po p
where xmlexists('/purchaseOrder/items/item[productName contains text {$x}]'
    passing p.object_value, 'Lapis' as "x");
```

Full Text Index Guideline 2: Pragma ora:no_schema: Using XML Schema-Based Data with XQuery Full Text

Oracle recommends in general that you use non XML Schema-based XMLType data stored as binary XML when you use XQuery Full Text and an XML full-text index. You can, however, use XML Schema-based XMLType data, provided you understand how it is used.

By default, when an XML full-text index is used to evaluate XML Schema-based data, compile-time error ORA-18177 is raised. This is because the full-text indexing itself makes no use of the associated XML schema: it is not type-aware. It treats all of the text that it applies to as untyped. This error is raised even if you type-cast data appropriately and thus do not depend on the XML schema to cast types implicitly. Example 61 illustrates this.

Example 61: XQuery Full Text Query with XML Schema-Based Data: Error ORA-18177

```
SELECT XMLQuery('/PurchaseOrder/LineItems/LineItem'
                PASSING OBJECT_VALUE RETURNING CONTENT)
FROM oe.purchaseorder
WHERE XMLEExists('/PurchaseOrder
                 [LineItems/LineItem/@ItemNumber > xs:integer("20")
                 and Actions/Action/User contains text "KPARTNER"]'
                 PASSING OBJECT_VALUE);
FROM oe.purchaseorder
*
```

ERROR at line 3:

```
ORA-18177: XQuery full text expression '/PurchaseOrder
[LineItems/LineItem/@ItemNumber > xs:integer("20")
and Actions/Action/User contains text "KPARTNER"]' cannot be evaluated using XML
text index
```

The error raised draws this to your attention, in case you might be expecting a full-text condition in your query to depend on XML Schema types and typed operations.

In order to use a condition that depends on types you must explicitly cast the relevant XQuery expressions to the appropriate types. Do not expect Oracle XML DB to use the XML schema to perform implicit type casting. Failure to type-cast appropriately can lead to results that you might not expect.

Example 62 shows a query of XML Schema-based data, with explicit type-casting to ensure that the proper condition is evaluated.

Example 62: Using XQuery Pragma ora:no_schema with XML Schema-Based Data

```
SELECT XMLQuery('/PurchaseOrder/LineItems/LineItem'
```

```

        PASSING OBJECT_VALUE RETURNING CONTENT)

FROM oe.purchaseorder

WHERE XMLEExists(' (# ora:no_schema #)

        {/PurchaseOrder

        [LineItems/LineItem/@ItemNumber > xs:integer("20")

        and Actions/Action/User contains text "KPARTNER"]}')

        PASSING OBJECT_VALUE);

```

However, most uses of XQuery Full Text expressions, even with XML Schema-based data, do not involve data that is typed. Just remember that if you do use a condition that makes use of typed data then you need to cast to the proper type.

In sum, if you are sure that your query does not involve typed data, or if you judge that it is all right to treat particular typed data as if it were untyped, or if you explicitly type-cast any data that needs to be typed, then you can use Oracle XQuery pragma `ora:no_schema` in your query to inhibit raising the error and allow evaluation of the query using an XML full-text index.

Full Text Index Guideline 3: Pragma `ora:use_xmltext_idx`: Forcing the Use of an XML Full-Text Index

A given query involving XML data can be evaluated in various ways, depending on the existence of different indexes and other factors. Sometimes the default evaluation method is not the most performant and it would be more efficient to force the use of an existing XML full-text index. You can use XQuery pragma `ora:use_xmltext_idx` to do this.

For example, a WHERE clause might include two XMLEExists expressions, only one of which involves an XQuery full-text condition, and you might have an XMLIndex index that applies to the XMLEExists expression that has no full-text condition. With such a query it is typically more efficient to use an XML full-text index to evaluate the entire WHERE clause.

Even in some cases where there is no full-text condition in the query, the use of an XML full-text index can provide the most efficient query evaluation.

The query in Example 63 illustrates the use of pragma `ora:use_xmltext_idx`. Only the first of the XMLEExists clauses uses a full-text condition. Because of the pragma, the full-text index (`po_ctx_idx`, created in Example 55) is used for both XMLEExists clauses.

Example 63: Full-Text Query with XQuery Pragma `ora:use_xmltext_idx`

```

SELECT XMLQuery('/PurchaseOrder/LineItems/LineItem'

        PASSING OBJECT_VALUE RETURNING CONTENT)

FROM po_binxml

```

```
WHERE XMLEExists('/PurchaseOrder/LineItems/LineItem
                [Description contains text "Picnic"]' PASSING OBJECT_VALUE)
AND XMLEExists('(# ora:use_xmltext_idx #) {/PurchaseOrder[User="SBELL"]}'
                PASSING OBJECT_VALUE);
```

Structured (Object Relational) storage

For structured storage, XQuery optimization is done by rewriting the different elements to the underlying relational columns, and by rewriting the collection access to underlying relational collection tables. XPath rewrite for object-relational storage means that XQuery with XPath expression is rewritten to a SQL query block or expression on the underlying relational tables or columns. These underlying tables can include out-of-line tables. This section presents some guidelines for using execution plans to do the following, for queries that use XPath expressions:

- Analyze query execution, to determine whether XPath rewrite occurs.
- Optimize query execution, by using secondary indexes.

Each guideline is listed below as “Structured Storage Guideline”. Use these guidelines together, taking all that apply into consideration.

Note: It is assumed that the reader of this section has already read the whitepaper titled “Ease of use packages for XMLType Structured Storage” and is familiar with the DBMS_XMLSCHEMA_ANNOTATE and DBMS_XMLSCHEMA_MANAGE packages.

Structured Storage Guideline 1: Make the SQL types and XML types correspond

Please refer to Table 2 in XQuery Guideline 7: “Using proper XQuery and SQL Typing” to make sure SQL type is selected properly for schema type when registering the XML schema for structured storage. Since structured storage is only available for schema based XML, xdb:sqlType should be set to NUMBER to make “Example 13: using xs:decimal() type exact numeric comparison” query efficient and xdb:sqlType should be set to BINARY_DOUBLE to make “Example 14: Using xs:double() type approximate numeric comparison” query efficient.

Structured Storage Guideline 2: Look for underlying tables and columns versus XML functions in execution plans

The execution plan of a query that has been rewritten refers to the object-relational tables and columns that underlie the queried XMLType data.

The names of the underlying tables can be meaningful to you, if they are derived from XML element or attribute names or if the governing XML schema explicitly names them by using annotation `xdb:defaultTable`. Otherwise, these names are system-generated and have no obvious meaning; in particular, they do not reflect the corresponding XML element or attribute names. Also, some system-generated columns are hidden; you do not see them if you use the SQL `describe` command. They nevertheless show up in execution plans.

The plan of a query that has not been rewritten shows only the base table names, and it typically refers to user-level XML functions, such as `XMLExists`. Look for this difference to determine

whether a query has been optimized. The XML function name shown in an execution plan is actually the internal name (for example, `XMLEXISTS2`), which is sometimes slightly different from the user-level name.

The example below shows the kind of execution plan output that is generated when Oracle XML DB cannot perform XPath rewrite. The plan here is for a query that uses SQL function `XMLEXISTS`. The corresponding internal function `XMLEXISTS2` appears in the plan output, indicating that the query is not rewritten.

Example 47: Execution Plan Generated When XPath Rewrite Does Not Occur

Predicate Information (identified by operation id):

```
1 - filter (XMLEXISTS2('$p/PurchaseOrder[User="SBELL]') PASSING BY VALUE
SYS_MAKEXML('6168797E040',4215,"PO"."XMLEXTRA","PO"."XMLDATA") AS "p")=1)
```

In this situation, Oracle XML DB constructs a pre-filtered result set based on any other conditions specified in the query `WHERE` clause. It then filters the rows in this potential result set to determine which rows belong in the result set. The filtering is performed by constructing a DOM on each document and performing a functional evaluation (using the methods defined by the DOM API) to determine whether or not each document is a member of the result set.

The example below shows a query, together with its execution plan, which shows that the query has been optimized. The predicates for `CostCenter` and `User` are rewritten to the underlying relational column filters.

Example 48: Optimization of XMLQuery with Schema-Based XMLType Data

```
SELECT XMLQuery('for $i in /PurchaseOrder
  where $i/CostCenter eq "A10"
  and $i/User eq "SMCCAIN"
  return <A10po pono="{ $i/Reference }"/>'
  PASSING OBJECT_VALUE
  RETURNING CONTENT)
FROM purchaseorder;
```

PLAN_TABLE_OUTPUT

Id	Operation	Name	Rows	Bytes	Cost
0	SELECT STATEMENT		1	530	5
1	SORT AGGREGATE		1		
* 2	FILTER				
3	FAST DUAL		1		2
* 4	TABLE ACCESS FULL	PURCHASEORDER	1	530	5

Predicate Information (identified by operation id):

```
2 - filter (:B1='SMCCAIN' AND :B2='A10')
4 - filter (SYS_CHECKACL("ACLOID", "OWNERID", xmltype('<privilege
xmlns="http://xmlns.oracle.com/xdb/acl.xsd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://xmlns.oracle.com/xdb/acl.xsd
http://xmlns.oracle.com/xdb/acl.xsd
DAV:http://xmlns.oracle.com/xdb/dav.xsd">
<read-properties/><read-contents/></privilege>'))=1)
```

Structured Storage Guideline 3: Name default tables and nested tables, so you recognize them in execution plans

When designing an XML schema, use annotation `xdb:defaultTable` to name the underlying tables that correspond to elements that you select in queries where performance is important. This lets you easily recognize them in an execution plan, indicating by their presence or absence whether the query has been rewritten.

When creating the XMLType tables using the CREATE TABLE statement, the nested tables can be renamed by using the “VARRAY ... STORE AS TABLE” clause.

If the nested tables need to be renamed after the table has been created by schema registration or CREATE TABLE statement, use the `DBMS_XMLSCHEMA_MANAGE.RENAMECOLLECTIONTABLE` procedure.

Structured Storage Guideline 4: Create relevant indexes

You can effectively index XML data that is stored object-rationally (structured storage) by creating B-tree indexes on the underlying database columns that correspond to XML nodes. Creating B-tree indexes will give near-relational performance for queries on structured storage, which gives structured storage an upper hand compared to Oracle’s competitors’ storage solutions, and even other XMLType storage solutions in Oracle XML DB. Sometimes, however, a query is better served by creating function-based index. This section talks about when and how to create B-tree and function-based indexes for structured storage. Each guideline is listed below as “Structured Storage Indexing Guideline”.

Structured Storage Indexing Guideline 5: Create an index on a column targeted by a predicate

Rewrite optimizations of XPath predicate, WHERE predicate of XQuery, WHERE clause of SQL statement may all result in searching values in columns of the XML storage or index tables. When this happens, you can sometimes improve performance by creating an index on the column that is targeted by the SQL predicate, or by creating an index on a function application to that column. There are several ways you can create an index on the predicate. Many of these are listed below:

Scenario 1: Using the XPath to create B-tree index

If the data to be indexed is a singleton, that is, if it can occur only once in any XML instance document, and you know the XPath to index, then you can use a shortcut of ostensibly creating a function-based index, where the expression defining the index is a functional application, with an XPath-expression argument that targets the singleton data. A shortcut is defined for XMLCast applied to XMLQuery, which is equivalent to the shortcut defined for extractValue. In many cases, Oracle XML DB then automatically creates appropriate B-Tree index on the underlying

object-relational tables or columns; it does not create a function-based index on the targeted XMLType data as the CREATE INDEX statement would suggest.

The example below shows a CREATE INDEX statement that ostensibly tries to create a function-based index using XMLCast applied to XMLQuery, targeting the text content of singleton element Reference. The content of this element is only text, so targeting the element is the same as targeting its text node using node test text().

Example 49: CREATE INDEX with XMLCAST and XMLQUERY on a Singleton Element

```
CREATE INDEX po_reference_ix ON purchaseorder
  (XMLCast(XMLQuery ('$p/PurchaseOrder/Reference'
    PASSING po.OBJECT_VALUE AS "p" RETURNING CONTENT)
    AS VARCHAR2(128)));
```

In this example, Oracle XML DB rewrites the CREATE INDEX statements to create B-tree index on the underlying scalar data.

In some cases when you use either of these shortcuts, the CREATE INDEX statement is not able to create an index on the underlying scalar data as described, and it instead actually does create a function-based index on the referenced XMLType data. This is so, even if the value of the index might be a scalar. If this happens, you should drop the index and follow one of the other scenarios to create the B-tree index.

Scenario 2: Converting XPath to table / column name and creating the index

If you are using the XPath based access, you can use the DBMS_XMLSCHEMA_MANAGE.XPATH2TABCOLMAPPING procedure to get the table and column name to index as shown in the example below. This will take into account the nested tables and out of line tables to give you the final table / column name to index.

Example 50: Creating an Index on a Column Targeted by a Predicate

```
-- manually creating index on a given XPath expression.
select XDB.DBMS_MANAGE_XMLSTORAGE.XPath2TabColMapping(
  'PURCHASEORDER_TAB',NULL, '/ipo:purchaseOrder/ Reference ',
  '''http://www.example.com/IPO''' as "ipo") from dual;

-- Result should look like this (SYS_NC... could have different value)
-- <Result>
--   <Mapping TableName="PURCHASEORDER_TAB" ColumnName="SYS_NC00009$"/>
-- </Result>

-- now we could create an index or constraint by manually extracting the
required information
-- create index shipto_idx on PURCHASEORDER_TAB (SYS_NC00009$);
```

Note that “” and “” both create index on the same underlying relational column. In other words, they are 2 ways of performing the same operation.

Scenario 3: When XPath is not known

In some cases, the XPath is not easily determined. For example, you may have an involved XQuery that internally rewrites into a storage column. In such cases, you can analyze the explain plan output (as shown in “”) to determine which columns to create indexes on (“”).

Example 51: Analyzing an Execution Plan to Determine a Column to Index

Predicate Information (identified by operation id):

```
1 - filter(CAST("PURCHASEORDER"."SYS_NC00021$" AS VARCHAR2(128))
          ='Sarah J. Bell' AND
          SYS_CHECKACL("ACLOID", "OWNERID", xmltype('<privilege
            xmlns="http://xmlns.oracle.com/xdb/acl.xsd"
            xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
            xsi:schemaLocation="http://xmlns.oracle.com/xdb/acl.xsd
            http://xmlns.oracle.com/xdb/acl.xsd
            DAV:http://xmlns.oracle.com/xdb/dav.xsd
            "><read-properties/><read-contents/></privilege>'))=1)
```

The predicate information indicates that the expression `XMLCast(XMLQuery...)` is rewritten to an application of SQL function `cast` to the underlying relational column that stores the requestor information for the purchase order, `SYS_NC00021$`. This column name is system-generated.

The execution plan refers to this system-generated name, in spite of the fact that the governing XML schema uses annotation `SQLName` to name this column `REQUESTOR`.

Because these two names (user-defined and system-generated) refer to the same column, you can create a B-tree index on this column using either name. The example below shows these two equivalent ways to create the B-tree index on the predicate-targeted column.

Example 52: Creating an Index on a Column Targeted by a Predicate using several different ways

```
CREATE INDEX requestor_index ON purchaseorder ("SYS_NC00021$");
CREATE INDEX requestor_index ON purchaseorder ("XMLDATA"."REQUESTOR");
```

Scenario 4: Creating function based index

For the plan shown in “”, it makes sense to create a function-based index, using a functional expression that matches the one in the rewritten query. The example below illustrates this.

Example 53: Creating a Function-Based Index for a Column Targeted by a Predicate, and execution plan that indicates that the index is picked up

```
CREATE INDEX requestor_index ON purchaseorder
  (cast("XMLDATA"."REQUESTOR" AS VARCHAR2(128)));
```

Id	Operation	Name	Rows	Bytes
0	SELECT STATEMENT		1	524
* 1	TABLE ACCESS BY INDEX ROWID	PURCHASEORDER	1	524
* 2	INDEX RANGE SCAN	REQUESTOR_INDEX	1	

Predicate Information (identified by operation id):

```
1 - filter(SYS_CHECKACL("ACLOID", "OWNERID", xmltype('<privilege
```

```

xmlns="http://xmlns.oracle.com/xdb/acl.xsd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://xmlns.oracle.com/xdb/acl.xsd
                    http://xmlns.oracle.com/xdb/acl.xsd
                    http://xmlns.oracle.com/xdb/acl.xsd"
DAV:http://xmlns.oracle.com/xdb/dav.xsd">
  <read-properties/><read-contents/></privilege>')=1)

```

```
2 - access(CAST("SYS_NC00021$" AS VARCHAR2(128))='Sarah J. Bell')
```

In the particular case of this query, the original functional expression applies XMLCast to XMLQuery to target a singleton element, Requestor. This is a special case, where you can, as a shortcut, use such a functional expression directly in the CREATE INDEX statement, and that statement is rewritten to create an index on the underlying scalar data. In other words, the example below which targets an XPath expression, has the same effect as the example above, which targets the corresponding object-relational column.

Example 54: Creating a Function-Based Index for a Column Targeted by a Predicate

```

CREATE INDEX requestor_index
  ON purchaseorder po
  (XMLCast(XMLQuery('$p/PurchaseOrder/Requestor'
                  PASSING po.OBJECT_VALUE AS "p" RETURNING CONTENT)
          AS VARCHAR2(128)));

```

Structured Storage Indexing Guideline 6: Create indexes on ordered collection tables

If a collection is stored as an ordered collection table or an XMLType instance, then you can directly access members of the collection. Each member of the collection becomes a row in a table, so you can access it directly with SQL.

You can often improve performance by indexing such collection members. You do this by creating a composite index on (a) the object attribute that corresponds to the collection XML element or its attribute and (b) pseudocolumn NESTED_TABLE_ID.

The example below shows the execution plan for a query to find the Reference elements in documents that contain an order for given part ID. The collection of LineItem elements is stored as rows in the ordered collection table lineitem_table. Instead of using table purchaseorder from sample database schema HR, for illustration we manually create a new purchaseorder table (in a different database schema) with the same properties and same data, but having OCTs with user-friendly names. This can be done by using the “VARRAY ... STORE AS TABLE” clause in the CREATE TABLE statement.

Example 55: Execution Plan for a Selection of Collection Elements

```

SELECT XMLCast(XMLQuery('$p/PurchaseOrder/Reference'
                      PASSING OBJECT_VALUE AS "p" RETURNING CONTENT)
          AS VARCHAR2(4000)) "Reference"
FROM purchaseorder

```

```

WHERE
XMLExists('$p/PurchaseOrder/LineItems/LineItem/Part[@Id="717951002372"]'
          PASSING OBJECT_VALUE AS "p");

```

Id	Operation	Name	Rows
0	SELECT STATEMENT		1
1	NESTED LOOPS		
2	NESTED LOOPS		1
3	SORT UNIQUE		1
* 4	TABLE ACCESS FULL	LINEITEM_TABLE	1
* 5	INDEX UNIQUE SCAN	LINEITEM_TABLE_MEMBERS	1
6	TABLE ACCESS BY INDEX ROWID	PURCHASEORDER	1

Predicate Information (identified by operation id):

```

4 - filter("SYS_NC00009$" IS NOT NULL AND
"SYS_NC00011$"='717951002372')
5 - access("NESTED_TABLE_ID"="PURCHASEORDER"."SYS_NC0003400035$")

```

The execution plan shows a full scan of ordered collection table `lineitem_table`. This could be acceptable if there were only a few hundred documents in the `purchaseorder` table, but it would be unacceptable if there were thousands or millions of documents in the table.

To improve the performance of such a query, you can create an index that provides direct access to pseudocolumn `NESTED_TABLE_ID`, given the value of attribute `Id`. Unfortunately, Oracle XML DB does not allow indexes on collections to be created using XPath expressions directly.

The best way to create an index on the nested table is by using the `DBMS_XMLSCHEMA_MANAGE.XPATH2TABCOLMAPPING` procedure to get the table and column name to index, and then create a composite index of this column with pseudocolumn `NESTED_TABLE_ID`. This will take into account the nested tables and out of line tables to give you the final table and column names to index. The example below shows how to automate this process.

Example 56: Creating an Index on a nested Table Column Targeted by a Predicate

```

-- pl/sql program to automatically create one index for each mapping
-- of a given XPath expression.
DECLARE
  index_name VARCHAR2(32) := 'po_idx';
  tab_name   VARCHAR2(32);
  col_name   VARCHAR2(32);
  CURSOR idx_cursor IS
    SELECT "tab_name", "col_name"
    FROM XMLTable('/Result/Mapping'
                 passing (
                   SELECT XDB.DBMS_MANAGE_XMLSTORAGE.XPath2TabColMapping(

```

```

        'PURCHASEORDER_TAB',
        NULL,
        '/ipo:purchaseOrder/items/item/Part/Id',
        '''http://www.example.com/IPO''' as "ipo")
    FROM dual)
    COLUMNS "tab_name" VARCHAR2(32) PATH '/Mapping/@TableName',
            "col_name" VARCHAR2(32) PATH '/Mapping/@ColumnName');
BEGIN
    FOR entry IN idx_cursor LOOP
        EXECUTE IMMEDIATE 'CREATE INDEX ' || index_name || '_'
            || entry."col_name" || ' ON "' || entry."tab_name"
            ||'" ( "' || entry."col_name" || ', NESTED_TABLE_ID)';
    END LOOP;
END;
/

```

The example above is written in a generic fashion with a PLSQL cursor created to handle multiple XPath's. If you are certain that your XPath2TabColName function will only return one XPath, then you can use SQL similar to “” by just plugging in your particular XPath.

Querying resources in XML DB repository

Oracle XML DB provides two SQL functions, `equals_path` and `under_path`, that can be used to perform folder-restricted queries. Such queries limit SQL statements that operate on the `RESOURCE_VIEW` or `PATH_VIEW` to documents that are at a particular location in Oracle XML DB folder hierarchy. Do not use `ANY_PATH` or `PATH` in where clause.

Appendix A: Semantic differences between the deprecated mainly XPath 1.0 based functions and standard SQL/XML XQuery based functions

There are some important differences between the deprecated and the XQuery based syntax, which are listed below to make the migration easier for the users.

In the de-supported `extract()`, `existsNode()`, `table(xmlsequence())`, `extractValue()`, only XPath 1.0 can be used in the path expression. The SQL/XML standard operators `XMLQuery()`, `XMLExists()`, `XMLTable`, `XMLCast()` use XQuery 1.0 in the query expression. Other than this important difference, there are several other non-standard behavior in the de-supported operators that users must pay special attention when migrating to use the standard based operators.

- Schema based datatype comparison: When de-supported operators are applied to schema based XMLType column (object relational based structured storage and binary XML storage), the schema based datatype comparison semantics is applied, for example, comparing non-string type with string results in casting and datatype specific comparison. However, in the standard operators, **XQuery date type casting functions must be used**. Otherwise an error will be raised. See XQuery Guideline 7.

Example:

Assume `@podate` is `xs:date` type and `@poid` is `xs:integer` type and `purchaseOrder` is an XMLType table storing schema based `purchaseOrder` XML document instances.

De-supported syntax:

```
Select 1 from purchaseOrder p
where existsNode(value(p), '/PurchaseOrder[@podate > "1998-09-02"]') = 1
```

Standard based syntax:

```
Select 1 from purchaseOrder p where xmlexists( 'declare namespace po =
http://www.po.com;/PurchaseOrder[@podate >xs:date("1998-09-02"]' passing
value(p))
```

The following query raises type errors

```
Select 1 from purchaseOrder p where xmlexists( 'declare namespace po =
http://www.po.com;/PurchaseOrder[@podate > "1998-09-02"]' passing value(p))
```

De-supported syntax:

```
Select 1 from purchaseOrder p
where existsNode(value(p), '/PurchaseOrder[@poid = "3456"]') = 1
```

Standard based syntax:

```
Select 1 from purchaseOrder p where xmlexists('declare namespace po =
http://www.po.com;/PurchaseOrder[@poid = 3456]' passing value(p))
```

- Namespace patching: As the example shown above, **the namespace declaration must be specified unless the XML document has no namespace** whereas in the de-supported syntax, the namespace might be patched even if it is NOT specified as the third parameter of the operator.
- existsNode returns 0 or 1 while XMLEExists returns Boolean, so you can use new syntax in the SQL WHERE clause directly. To use it in the SELECT list, please refer to “Example 3: Using XMLEExists() with CASE Expression in select list”.
- Bind variable: There is no need to use string concatenation operator || to construct XPath string to embed bind variable as in the de-supported syntax. Instead, use PASSING clauses to pass bind variables to XQuery based functions.

Example:

De-supported syntax:

```
Select value(p) from purchaseOrder p
where existsNode(value(p), '/PurchaseOrder[@podate >' || ':1: ?'] = 1
```

Standard syntax:

```
Select value(p) from purchaseOrder p
where xmlexists('declare namespace po = http://www.po.com;/PurchaseOrder[@podate
> xs:date($bindvar)]' passing value(p), :1 as “bindvar”)
```

- ora:instanceof() and ora:instanceof-only() are only usable in the XPath of the de-supported syntax. Use XQuery ‘instance of P’ expression and ‘@xsi:type =’ respectively in the standard syntax.

Example: ora:instanceof()

De-supported syntax:

```
select extract(value(r),'/N2:R1[ora:instanceof(.,"N1:superType1")]',
'xmlns:N1="http://www.oracle.com/xdb/N1"
xmlns:N2="http://www.oracle.com/xdb/N2"
xmlns:ora="http://xmlns.oracle.com/xdb")' from R1 r;
```

Standard syntax:

```
select XMLQuery('declare namespace N1="http://www.oracle.com/xdb/N1";
declare namespace N2="http://www.oracle.com/xdb/N2";
/N2:R1[ instance of element(N2:R1, N1:superType1)]'
passing object_value returning content) from R1 r ;
```

Example: ora:instanceof-only()

De-supported syntax:

```
select extract(value(r),'/N2:R1[ora:instanceof-only(.,"N1:superType1")]',
'xmlns:N1="http://www.oracle.com/xdb/N1"')
```

```
xmlns:N2="http://www.oracle.com/xdb/N2"
xmlns:ora="http://xmlns.oracle.com/xdb") from R1 r;
```

Standard syntax:

```
select XMLQuery('declare namespace N1="http://www.oracle.com/xdb/N1";
declare namespace N2="http://www.oracle.com/xdb/N2";
/N2:R1[@xsi:type="N1:superType1"]' passing object_value returning content)
from R1 r;
```

Notice that xsi:type predicate is also supported in XPath, i.e., the following query works the same as the two above:

```
select extract(value(r),'/N2:R1[@xsi:type="N1:superType1"]',
'xmlns:N1="http://www.oracle.com/xdb/N1"
xmlns:N2="http://www.oracle.com/xdb/N2"
xmlns:ora="http://xmlns.oracle.com/xdb") from R1 r;
```

ora:upper(), ora:lower(), ora:to_number(), ora:to_date() are only usable in the XPath of the de-supported syntax. Use corresponding XQuery F&O functions fn:upper-case(), fn:lower-case(), xs:decimal(), xs:date() respectively in the standard syntax.

Example: DBMS_XMLGEN:

De-support syntax:

```
SELECT sys_XMLGen(km_t(kid,
kname,
knum,
CAST(MULTISET (SELECT kid, kdid, kdname
FROM ktest_d d
WHERE d.kid = m.kid) AS kdlist_t))),getclobval() AS detail
FROM ktest_m m;
```

Standard syntax:

```
select XMLSERIALIZE
(
document
XMLELEMENT
(
"KD_LIST",
XMLAGG
(
(
```

```
SELECT XMLAGG
(
  XMLELEMENT
  (
    "KD_T",
    XMLELEMENT("KID",KID),
    XMLELEMENT("KDID",KDID),
    XMLELEMENT("KDNAME",KDNAME)
  )
)
from KTEST_D d
where d.kid = m.kid
)
)
as clob indent size=2
)
from KTEST_M m;
```



White Paper Title
January 2013
Author: Oracle XML DB Team
Contributing Authors: [OPTIONAL]

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



| Oracle is committed to developing practices and products that help protect the environment

Copyright © 2013, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.