Oracle® Database JavaScript Developer's Guide





Oracle Database JavaScript Developer's Guide, Release 23c

F56885-03

Copyright © 2022, 2023, Oracle and/or its affiliates.

Primary Author: Sarah Hirschfeld

Contributors: M. Bach, L. Braun-Lohrer, H. Kasture, A. Ulrich, G. Venzl, M. Brantner, L. Daynes, H. Guiroux, A. Schubert, A. Burlison, M. Keppner, A. Kashuba, N. Sheikh, A.A. Baha, D. Adams

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, and MySQL are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

Introdu	uction to Oracle Database Multilingual Engine for JavaScript	
1.1 Th	e Need for a Multilingual Engine	1-2
1.2 Ov	verview of JavaScript	1-3
1.3 Ov	rerview of Multilingual Engine for JavaScript	1-3
1.3.1	JavaScript Implementation Details	1-4
1.3.2	2 Invoking JavaScript in the Database	1-5
1.3.3	B Introduction to Dynamic Execution	1-5
1.3.4	Introduction to MLE Module Calls	1-6
1.3.5	About MLE Execution Contexts	1-8
1.4 Int	roduction to Debugging JavaScript Code	1-9
	avaScript Modules and Environments	
	ing JavaScript Modules in MLE	2-1
2.1.1		2-3
	2.1.1.1 Naming JavaScript Modules	2-3
	2.1.1.2 Creating JavaScript Modules in the Database	2-4
	2.1.1.3 Storing JavaScript Code in Databases Using Single-Byte Character Sets	2-5
	2.1.1.4 Code Analysis	2-!
2.1.2		2-6
2.1.3	·	2-8
2.1.4 2.1.5		2-8
2.1.6	·	2-9 2-10
2.1.0	·	2-10
2.1.7	·	2-10
	B Dictionary Views Related to MLE JavaScript Modules 2.1.8.1 USER_SOURCE	2-1.
	2.1.8.2 USER_MLE_MODULES	2-1.
	ecifying Environments for MLE Modules	2-12
2.2.3p		2-17
	2.2.1.1 Naming MLE Environments	2-13
	2.2.1.2 Creating an Empty MLE Environment	2-12
	2.2.1.3 Creating an Environment as a Clone of an Existing Environment	2-15
	1.2.1.0 Stoking an Environment as a Sione of an Existing Environment	T



2.2.1.4 Using MLE Environments for Import Resolution	2-15
2.2.1.5 Providing Language Options	2-17
2.2.2 Dropping MLE Environments	2-18
2.2.3 Modifying MLE Environments	2-19
2.2.3.1 Altering Language Options	2-19
2.2.3.2 Modifying Module Imports	2-19
2.2.4 Dictionary Views Related to MLE JavaScript Environments	2-20
2.2.4.1 USER_MLE_ENVS	2-20
2.2.4.2 USER_MLE_ENV_IMPORTS	2-20
Overview of Dynamic MLE Execution	
3.1 About Dynamic JavaScript Execution	3-1
3.2 Dynamic Execution Workflow	3-2
3.2.1 Providing JavaScript Code Inline	3-2
3.2.2 Loading JavaScript Code from Files	3-3
3.3 Returning the Result of the Last Execution	3-6
Overview of Importing MLE JavaScript Modules	
4.1 JavaScript Module Hierarchies	4-2
4.1.1 Resolving Import Names Using MLE Environments	4-2
4.2 Export Functionality	4-3
4.2.1 Named Exports	4-3
4.2.2 Default Exports	4-4
4.2.3 Private Identifiers	4-5
4.3 Import Functionality	4-5
4.3.1 Module Objects	4-6
4.3.2 Named Imports	4-6
4.3.3 Default Imports	4-7
MLE JavaScript Functions	
5.1 Call Specifications for Functions	5-1
5.1.1 Creating a Call Specification for an MLE Module	5-2
5.1.1.1 Components of an MLE Call Specification	5-4
5.1.1.2 MLE Module Clause	5-5
5.1.1.3 ENV Clause	5-5
5.1.1.4 SIGNATURE Clause	5-6
5.1.2 Creating an Inlined MLE Call Specification	5-7
5.1.2.1 Components of an Inlined MLE Call Specification	5-9
5.1.2.2 Accessing Built-in Modules Using JavaScript Global Variables	5-10
· · · · · · · · · · · · · · · · · · ·	





7.0 116		
7.3 Usi	ng SODA for In-Database JavaScript	7-3
7.3.1	Getting Started with SODA for In-Database JavaScript	7-5
7.3.2	Creating a Document Collection with SODA for In-Database JavaScript	7-7
7.3.3	Opening an Existing Document Collection with SODA for In-Database JavaScript	7-9
7.3.4	Checking Whether a Given Collection Exists with SODA for In-Database JavaScript	7-10
7.3.5	Discovering Existing Collections with SODA for In-Database JavaScript	7-10
7.3.6	Dropping a Document Collection with SODA for In-Database JavaScript	7-11
7.3.7	Creating Documents with SODA for In-Database JavaScript	7-12
7.3.8	Inserting Documents into Collections with SODA for In-Database JavaScript	7-15
7.3.9	Saving Documents into Collections with SODA for In-Database JavaScript	7-17
7.3.10	SODA for In-Database JavaScript Read and Write Operations	7-19
7.3.11	Finding Documents in Collections with SODA for In-Database JavaScript	7-20
7.3.12	Replacing Documents in a Collection with SODA for In-Database JavaScript	7-26
7.3.13	Removing Documents from a Collection with SODA for In-Database JavaScript	7-27
	Indexing the Documents in a Collection with SODA for In-Database	
7.3.14	JavaScript	7-29
7.3.1 ⁴ 7.3.1 ⁵	JavaScript	7-29 7-31
	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript	
7.3.1	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript	7-31
7.3.19 7.3.10 7.3.1	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript	7-31 7-32
7.3.19 7.3.10 7.3.11 Post-E	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript xecution Debugging of MLE JavaScript Modules	7-31 7-32
7.3.19 7.3.10 7.3.11 Post-E	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Secifying Debugpoints	7-31 7-32 7-32
7.3.19 7.3.10 7.3.11 Post-E 8.1 Spe 8.1.1	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Debugpoints Debugpoint Locations	7-31 7-32 7-32 8-2 8-2
7.3.19 7.3.17 7.3.17 Post-E 8.1 Spe 8.1.1 8.1.2	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Debuggoint Locations Debuggoint Actions	7-31 7-32 7-32 8-2 8-2 8-2
7.3.19 7.3.17 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Debugpoint Locations Debugpoint Actions Debugpoint Conditions	7-31 7-32 7-32 8-2 8-2
7.3.19 7.3.17 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Debugpoint Locations Debugpoint Actions Debugpoint Conditions Database JavaScript Debugpoint Actions Debugpoint Conditions Debugpoint Debugpoints	7-31 7-32 7-32 8-2 8-2 8-2 8-4
7.3.19 7.3.10 7.3.11 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3 8.2 Mai	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Debugpoint Locations Debugpoint Actions Debugpoint Conditions Database JavaScript Debugpoint Conditions Debugging Security Considerations	7-31 7-32 7-32 8-2 8-2 8-2 8-4 8-4
7.3.19 7.3.10 7.3.11 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3 8.2 Mai 8.2.1 8.2.2	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Ecifying Debugpoints Debugpoint Locations Debugpoint Actions Debugpoint Conditions naging Debugpoints Debugging Security Considerations COLLECT DEBUG INFO Privilege for MLE Modules	7-31 7-32 7-32 8-2 8-2 8-2 8-4 8-4 8-6
7.3.19 7.3.10 7.3.11 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3 8.2 Mai 8.2.1 8.2.2 8.3 Ana	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Debugpoint Locations Debugpoint Actions Debugpoint Conditions Debugging Debugpoints Debugging Security Considerations COLLECT DEBUG INFO Privilege for MLE Modules Debug Output	7-31 7-32 7-32 8-2 8-2 8-4 8-4 8-6 8-7
7.3.19 7.3.10 7.3.11 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3 8.2 Mai 8.2.1 8.2.2	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Ecifying Debugpoints Debugpoint Locations Debugpoint Actions Debugpoint Conditions The provided Handle Security Considerations COLLECT DEBUG INFO Privilege for MLE Modules Edyzing Debug Output Textual Representation of Debug Output	7-31 7-32 7-32 8-2 8-2 8-4 8-4 8-6 8-7 8-7
7.3.19 7.3.10 7.3.11 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3 8.2 Mai 8.2.1 8.2.2 8.3 Ana 8.3.1 8.3.2	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Debuggoints Debugpoint Locations Debugpoint Actions Debugpoint Conditions Debugging Security Considerations COLLECT DEBUG INFO Privilege for MLE Modules Debug Output Textual Representation of Debug Output Analyzing Debug Output Using Developer Tools	7-31 7-32 7-32 8-2 8-2 8-4 8-4 8-6 8-7 8-7
7.3.19 7.3.10 7.3.11 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3 8.2 Mai 8.2.1 8.2.2 8.3 Ana 8.3.1 8.3.2	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Ecifying Debugpoints Debugpoint Locations Debugpoint Actions Debugpoint Conditions The provided Handle Security Considerations COLLECT DEBUG INFO Privilege for MLE Modules Edyzing Debug Output Textual Representation of Debug Output	7-31 7-32 7-32 7-32 8-2 8-2 8-4 8-4 8-6 8-7 8-7 8-7 8-7
7.3.19 7.3.10 7.3.11 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3 8.2 Mai 8.2.1 8.2.2 8.3 Ana 8.3.1 8.3.2 8.4 Erro	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript XECUTION Debugging of MLE JavaScript Modules Ecifying Debugpoints Debugpoint Locations Debugpoint Actions Debugpoint Conditions Debugging Security Considerations COLLECT DEBUG INFO Privilege for MLE Modules Alyzing Debug Output Textual Representation of Debug Output Analyzing Debug Output Using Developer Tools Or Handling in MLE Errors in Callouts	7-31 7-32 7-32 7-32 8-2 8-2 8-4 8-4 8-6 8-7 8-7 8-7 8-10 8-10
7.3.19 7.3.10 7.3.11 Post-E 8.1 Spe 8.1.1 8.1.2 8.1.3 8.2 Mai 8.2.1 8.2.2 8.3 Ana 8.3.1 8.3.2 8.4 Erro 8.4.1 8.4.2	JavaScript Getting a Data Guide for a Collection with SODA for In-Database JavaScript Handling Transactions with SODA for In-Database JavaScript Summary Example of Using SODA for In-Database JavaScript Execution Debugging of MLE JavaScript Modules Debugpoint Debugging of MLE JavaScript Modules Debugpoint Actions Debugpoint Actions Debugpoint Conditions Database JavaScript Modules Debugpoint Locations Debugpoint Actions Debugpoint Conditions Debugging Security Considerations COLLECT DEBUG INFO Privilege for MLE Modules Debugging Debug Output Textual Representation of Debug Output Analyzing Debug Output Using Developer Tools THANDLING THANDLING TOOLS TOOLS	7-31 7-32 7-32 7-32 8-2 8-2 8-4 8-4 8-6 8-7 8-7 8-7 8-10 8-10 8-13



9 MLE Security

9.1	Syst	em and Object Privileges Required for Working with JavaScript in MLE	9-1
	9.1.1	Necessary Privileges for the Execution of JavaScript Code	9-2
	9.1.2	Necessary Privileges for Using the NoSQL API	9-2
	9.1.3	Necessary Privileges for Creating MLE Schema Objects	9-3
	9.1.4	Necessary Privileges for Creating MLE Modules and Environments in ANY Schema	9-3
	9.1.5	Necessary Privileges for Post-Execution Debugging	9-4
9.2	Seci	urity Considerations for MLE	9-5
	9.2.1	Initialization Parameter	9-5
	9.2.2	Execution Contexts	9-6
	9.2.3	Runtime State Isolation	9-6
	9.2.4	Database Security Model	9-8
	9.2.5	Considerations for Using MLE Call Specifications and Modules from Different Schemas	9-9
	9.2.6	Auditing MLE Operations in Oracle Database	9-10
9.3	Java	Script Security Best Practices	9-11
	9.3.1	Using Bind Variables for Security and Performance	9-11
	9.3.2	Generic Database and PL/SQL Specific Security Considerations	9-13
	9.3.3	Supply Chain Security	9-13
	9.3.4	Software Bill of Material	9-14
	9.3.5	Using the Database to Store State	9-15
	9.3.6	Disabling Multilingual Runtime	9-16
9.4	MLE	Security Examples	9-17
	9.4.1	Business Logic Stored in MLE Modules	9-17
	9.4.2	Generic Data Processing Libraries	9-19
	9.4.3	Generic Libraries in Business Logic	9-20

A MLE Type Conversions

Index



List of Examples

2-1	Creating a JavaScript Module in the Database	2-4
2-2	Create a Call Specification for a Public Function	2-6
2-3	Public and Private Functions in a JavaScript Module	2-7
2-4	Providing JavaScript Source Code Using a BFILE	2-8
2-5	Specification of a VERSION string in CREATE MLE MODULE	2-9
2-6	Addition of JSON Metadata to the MLE Module	2-9
2-7	Drop an MLE Module	2-10
2-8	Drop an MLE Module Using IF EXISTS	2-10
2-9	Alter an MLE Module	2-10
2-10	Externalize JavaScript Module Source Code	2-11
2-11	Find MLE Modules Defined in a Schema	2-12
2-12	Map Identifier to JavaScript Module	2-16
2-13	Import Module Functionality	2-16
2-14	List Available MLE Environments Using USER_MLE_ENVS	2-20
2-15	List Module Import Information Using USER_MLE_ENV_IMPORTS	2-20
3-1	Using the Q-Quote Operator to Provide JavaScript Code Inline with PL/SQL	3-2
3-2	Loading JavaScript code from a BFILE with DBMS_LOB.LOADCLOBFROMFILE()	3-3
3-3	Loading JavaScript Code from a BFILE by Referencing an MLE Module from DBMS_MLE	3-5
3-4	Returning the Result of the Last Execution	3-6
4-1	Use an MLE Environment to Map an Import Name to a Module	4-2
4-2	Function Export using Named Exports	4-3
4-3	Function Export Using Export Keyword Inline	4-4
4-4	Export a Class Using a Default Export	4-4
4-5	Named Export of Single Function	4-5
4-6	Module Object Definition	4-6
4-7	Named Imports Using Specified Identifiers	4-6
4-8	Named Imports with Aliases	4-7
4-9	Default Import	4-7
4-10	Default Import with Built-in Module	4-8
5-1	Creating MLE Call Specifications	5-2
5-2	Simple Inlined MLE Call Specification	5-8
5-3	Inlined MLE Call Specification Returning JSON	5-8
5-4	Execution Context Dependencies	5-12
5-5	Show JavaScript Call Specification Metadata	5-15
5-6	OUT and IN OUT Parameters with JavaScript	5-16



6-1	Getting Started with the MLE JavaScript SQL Driver	6-3
6-2	Use Global Variables to Simplify SQL Execution	6-5
6-3	Selecting Data Using Direct Fetch: Arrays	6-7
6-4	Selecting Data Using Direct Fetch: Objects	6-8
6-5	Fetching Rows Using a ResultSet	6-9
6-6	Using the Iterable Protocol with ResultSets	6-10
6-7	Using the Iterator Protocol with ResultSets	6-11
6-8	Updating a Row Using the MLE JavaScript SQL Driver	6-12
6-9	Using Named Bind Variables	6-14
6-10	Using Positional Bind Variables	6-15
6-11	Using the RETURNING INTO Clause	6-15
6-12	Performing a Batch Operation	6-17
6-13	Calling PL/SQL from JavaScript	6-18
6-14	SQL Error Handling Inside a JavaScript Function	6-19
6-15	Error Handling Using JavaScript throw() Command	6-20
6-16	Inserting JSON Data into a Database Table	6-21
6-17	Inserting a CLOB into a Table	6-24
6-18	Read an LOB	6-25
6-19	Using JavaScript Native Data Types vs Using Wrapper Types	6-29
6-20	Overriding the Global oracledb.fetchAsPlsqlWrapper Property	6-29
7-1	SODA with MLE JavaScript General Workflow	7-5
7-2	Creating a Collection with Custom Metadata	7-8
7-3	Opening an Existing Document Collection	7-9
7-4	Fetching All Existing Collection Names	7-10
7-5	Filtering the List of Returned Collections	7-10
7-6	Dropping a Collection	7-12
7-7	Creating SODA Documents	7-13
7-8	Inserting a SODA Document into a Collection	7-15
7-9	Inserting an Array of Documents into a Collection	7-16
7-10	Saving a Document into a Collection	7-17
7-11	Finding a Document by Key	7-20
7-12	Looking up Documents Using Multiple Keys	7-21
7-13	Using a QBE to Filter Documents in a Collection	7-22
7-14	Using skip() and limit() in a Pagination Query	7-23
7-15	Specifying Document Versions	7-24
7-16	Counting the Number of Documents Found	7-25
7-17	Replacing a Document in a Collection and Returning the Result Document	7-26



7-18	Removing a Document from a Collection Using a Document Key	7-27
7-19	Removing JSON Documents from a Collection Using a Filter	7-28
7-20	Creating a B-Tree Index for a JSON Field with SODA for In-Database JavaScript	7-29
7-21	Creating a JSON Search Index with SODA for In-Database JavaScript	7-30
7-22	Dropping an Index with SODA for In-Database JavaScript	7-31
7-23	Generating a Data Guide for a Collection	7-32
7-24	Use SODA for In-Database JavaScript	7-33
8-1	JSON Template for Specifying Debugpoints	8-2
8-2	JSON Template for Specifying Watch Action	8-3
8-3	JSON Template for Specifying Snapshot Action	8-4
8-4	Watching a Variable in an MLE Module	8-4
8-5	Enabling Debugging of an MLE Module	8-5
8-6	Obtain Textual Representation of Debug Output	8-8
8-7	Throwing ORA-04161 Error and Querying the Stack Trace	8-11
8-8	Redirect stdout to CLOB and DBMS_OUTPUT for MLE Module	8-14
8-9	Redirect stdout to CLOB and DBMS_OUTPUT for Dynamic MLE	8-16
9-1	Runtime State Isolation Scenario	9-7
9-2	Using Bind Variables Rather than String Concatenation	9-12
9-3	Use DBMS_ASSERT to Verify Valid Input	9-12
9-4	Using Bind Variables Rather than String Concatenation	9-15
9-5	Use DBMS_ASSERT to Verify Valid Input	9-15
9-6	Business Logic Stored in MLE Modules	9-17
9-7	Generic Data Processing Libraries	9-19
9-8	Use Generic Libraries in Business Logic	9-20



List of Figures

5-1	MLE Call Specification Syntax	5-4
5-2	MLE Inlined Call Specification Syntax	5-9
7-1	SODA for In-Database JavaScript Basic Workflow	7-2



List of Tables

2-1	JavaScript Language Options	2-18
5-1	Components of an MLE Call Specification	5-4
5-2	Components of an Inlined MLE Call Specification	5-9
7-1	Overview of Nonterminal Methods for Read Operations	7-19
7-2	Overview of Terminal Methods for Read Operations	7-19
7-3	Overview of Terminal Methods for Write Operations	7-20
A-1	Supported Mappings from SQL and PL/SQL Types to JavaScript Types	A-2
A-2	Supported Mappings from JavaScript Types to SQL Types	A-3
A-3	Mapping from JSON Attribute Types and Values to JavaScript Types and Values	A-4
A-4	Mapping from JavaScript Types and Values to JSON Attributes and Values	A-5



1

Introduction to Oracle Database Multilingual Engine for JavaScript

Oracle Database supports a rich set of languages for writing user-defined functions and stored procedures, including PL/SQL, Java, and C. With Oracle Database Multilingual Engine (MLE), developers have the option to run JavaScript code through dynamic execution or with persistent MLE modules stored directly in the database.

The landscape of programming languages is rapidly evolving, with more developers choosing to use modern dynamic languages like JavaScript. Besides simpler syntax and support for modern language features, a key factor in the popularity of these languages is the existence of a rich module ecosystem. Developers often choose to use different languages to implement different parts of a project, based on the availability of suitable modules for the given task.

Whether or not a new language reaches widespread adoption frequently depends on community involvement. Once a language reaches some threshold of popularity, its ecosystem often starts expanding rapidly, attracting more and more developers. Many times, a rich set of features, libraries, and reusable code modules are created to support more widespread use.

The Oracle Database is renowned for its support of a rich ecosystem of programming languages. The most common programmatic server-side interface to the Oracle Database is PL/SQL. By using PL/SQL it is possible to keep business logic and data together, oftentimes offering significant improvements to efficiency in addition to providing a unified processing pattern for data, regardless of the client interface in use. With MLE, you can utilize PL/SQL to implement JavaScript modules, offering an additional avenue to interact directly with the database.



Oracle Database Development Guide for more information about the programming languages supported by the Oracle database.

Topics

- The Need for a Multilingual Engine
 The benefits of using MLE to process data within the database are described.
- Overview of JavaScript
 One of the most popular programming languages today, JavaScript runs on any machine with a JavaScript engine. Developers prefer JavaScript mainly for the ease of scripting to develop end-to-end applications and for fast execution.
- Overview of Multilingual Engine for JavaScript
 MLE allows you to run and store JavaScript directly in the Oracle Database.

Introduction to Debugging JavaScript Code
 MLE allows you to debug your JavaScript code by conveniently and efficiently collecting runtime state during program execution.

1.1 The Need for a Multilingual Engine

The benefits of using MLE to process data within the database are described.

When developers implement a *Smart-DB approach*, application logic and data coexist in the same database. Applying this strategy, the database is used as a full-fledged processing engine as opposed to simply a persistence layer or a simple REST API. Making use of the database for processing data where it lives can provide numerous advantages in the form of enhanced security, potential elimination of network round-trips, and better data quality thanks to the use of referential integrity.

The database's optimizer also benefits from this approach. Using referential integrity constraints allows it to know more about the data it's working with. Performance benefits can also be realized when using set-based SQL and oftentimes, database servers are more powerful than the machines serving the application's front-end, further speeding up processing time.

The Smart-DB approach requires you to be familiar with the programming languages offered by the database system to make the best use of the concept. The only other option is to use a client-side driver to extract data from the database to a middleware system or client machine for processing.

With the ever-increasing data volumes to be handled, especially for batch-processing, transferring large quantities of data from the database to a client can be problematic for the following reasons:

- The transfer of database information between servers is time consuming and can cause significant network overhead
- Latencies are often significantly increased; the cumulative effect can be very noticeable, especially for "chatty" applications
- Processing large data volumes in a middle-tier or client requires these environments to be equipped with large amounts of DRAM and storage, adding cost
- Data transfer between machines, especially in cloud environments, is often subject to regulatory control due to the inherent security risks and data protection requirements

Processing data *within* the database is a common strategy for mitigating against many of these problems.

With the introduction of Oracle Database Multilingual Engine (MLE), JavaScript is added to the database. The inclusion of JavaScript acknowledges the language's popularity and opens its extensive ecosystem for server-side database development.

With MLE, you can use idioms and tools available in JavaScript's ecosystem, as well as deploy and use modules from popular repositories such as Node Package Manager (NPM) right in the database. Furthermore, you can move between application tiers, providing more flexibility to teams dealing with varying workloads. The large pool of JavaScript talent can help staff existing and upcoming projects.



1.2 Overview of JavaScript

One of the most popular programming languages today, JavaScript runs on any machine with a JavaScript engine. Developers prefer JavaScript mainly for the ease of scripting to develop end-to-end applications and for fast execution.

JavaScript (JS) has come a long way since its inception as a browser-based solution for interactive web pages. While its popularity for front-end development remains strong, it has found its way into back-end development as well. For example, Node.js and Deno are very popular in that space.

At its core, JavaScript is an interpreted language with support for many modern programming styles. JavaScript is continually enhanced by a governing body known as ECMA International with new standards released annually.

JavaScript features both a functional as well as an object oriented interface. Despite the name, JavaScript is very different from Java, although its syntax intentionally mimics many constructs known in other popular languages. The learning curve is eased by providing a familiar looking syntax.

Soft factors such as a very large and active community as well as the language's rich set of libraries make it an attractive choice for development.

With the introduction of Oracle Database Multilingual Engine (MLE), it is possible to execute JavaScript directly in the Oracle database. Data-intensive applications can benefit from moving processing logic from the middle-tier to the database.



Developer.mozilla.org for more information about JavaScript

1.3 Overview of Multilingual Engine for JavaScript

MLE allows you to run and store JavaScript directly in the Oracle Database.

Using MLE enables users of the Oracle Database to run the following, written in JavaScript:

- Stored procedures
- Stored functions
- Code in a PL/SQL package namespace
- Anonymous, dynamic code snippets (in a way that is similar to DBMS SQL)

MLE is supported when connecting to the database using a dedicated server connection on Linux x86-64. Abstract data types, PL/SQL collections and record types, and REF CURSORS are not supported. Additionally, the initialization parameter MAX_STRING_SIZE must not have been changed from the default.



Note:

Shared server connections and those using Database Resident Connection Pool (DRCP) cannot make use of MLE.

Topics

- JavaScript Implementation Details
 The MLE implementation of JavaScript is compliant with ECMAScript 2022.
- Invoking JavaScript in the Database
 JavaScript can be invoked through dynamic execution or through call specifications, which either reference MLE modules or inlined JavaScript functions.
- Introduction to Dynamic Execution
 Anonymous JavaScript code snippets can be executed via the DBMS_MLE PL/SQL package.
- Introduction to MLE Module Calls
 It is possible to create JavaScript modules as schema objects that are stored persistently in the database.
- About MLE Execution Contexts
 An MLE execution context encapsulates all runtime state associated with the execution of JavaScript code in MLE.

1.3.1 JavaScript Implementation Details

The MLE implementation of JavaScript is compliant with ECMAScript 2022.

Adhering to the ECMA standard, the JavaScript implementation as found in MLE is consciously created as a *pure implementation*. Native JavaScript network and file I/O operations are not supported in the same way that they are in Node.js and Deno for security reasons. The use of network and file I/O is possible with MLE, however, you must employ PL/SQL APIs such as UTL_HTTP and UTL_FILE.

The WEB API, Fetch, is not available by default in the global space but can be enabled by importing mle-js-fetch.

Objects not included in the ECMA standard, including common objects used in frontend code such as the Window object, are also not available with MLE. Nevertheless, MLE does provide easy and efficient access to SQL, which is able to execute close to the data. Console output is passed to <code>DBMS_OUTPUT</code> by default but can be redirected and stored in a user provided CLOB if required.

Users require specific privileges before they can interact with MLE. These can broadly be classified into:

- Permission to use MLE and run JavaScript code
- Execute dynamic JavaScript in the database
- Create JavaScript modules and externalize them via PL/SQL code

The database engine throws an error if you lack sufficient privileges required for the use of JavaScript.





System and Object Privileges Required for Working with JavaScript in MLE for more information about privileges

1.3.2 Invoking JavaScript in the Database

JavaScript can be invoked through dynamic execution or through call specifications, which either reference MLE modules or inlined JavaScript functions.

Generally speaking, server-side JavaScript code can be invoked in two ways:

- Dynamically via the DBMS MLE package
- Using PL/SQL code referencing functions exported in JavaScript modules (so-called MLE module calls) or functions defined directly in the DDL

Regardless of which of the two methods is used, all JavaScript code runs in an execution context. Its purpose is to encapsulate all runtime state associated with the processing of JavaScript code. The MLE execution context corresponds to the ECMAScript execution context for JavaScript.



Ecma-international.org for more information about the ECMAScript execution context

1.3.3 Introduction to Dynamic Execution

Anonymous JavaScript code snippets can be executed via the DBMS MLE PL/SQL package.

The procedure <code>DBMS_MLE.eval()</code> is used to execute dynamic MLE snippets. The procedure takes the following arguments:

Argument Name	Туре	Optional?	
CONTEXT_HANDLE	RAW(16)	N	
LANGUAGE_ID	VARCHAR2(64)	N	
SOURCE	CLOB	N	
RESULT	CLOB	Y	
SOURCE NAME	VARCHAR2	Υ	

The argument <code>SOURCE_NAME</code> is optionally used to provide a name for the otherwise randomly-named <code>JavaScript</code> code block.

JavaScript code can be provided inline with PL/SQL as shown in the following code:

SET SERVEROUTPUT ON



Executing this example will result in the following being printed:

```
Hello World, this is DBMS_MLE
```

The code provided above demonstrates the following concepts of invoking JavaScript code dynamically:

- An execution context must be explicitly created
- JavaScript code is provided as a Character Large Object (CLOB) or VARCHAR2 variable
- The context must be explicitly evaluated

Both PL/SQL and JavaScript are present when you execute JavaScript dynamically. The code snippets provided are not reusable outside of their namespace. The output of the call to console.log is passed to DBMS OUTPUT for printing on the screen.

See Also:

- Overview of Dynamic MLE Execution for more details about dynamic execution with MLE
- Returning the Result of the Last Execution for more information about the RESULT argument of the procedure DBMS MLE.eval()

1.3.4 Introduction to MLE Module Calls

It is possible to create JavaScript modules as schema objects that are stored persistently in the database.



Once a JavaScript module has been defined, it can be used in SQL and PL/SQL as shown below:

```
CREATE OR REPLACE MLE MODULE helloWorld_module
LANGUAGE JAVASCRIPT AS
function helloWorld() {
    console.log('Hello World, this is a JS module');
}
export { helloWorld }
//
```

Before the exported JavaScript function can be invoked, a call specification must be defined. The code snippet below shows how to create a call specification for the JavaScript helloWorld() function in PL/SQL:

```
CREATE OR REPLACE PROCEDURE helloWorld_proc
AS MLE MODULE helloWorld_module
SIGNATURE 'helloWorld()';
/
```

The call specification, referred to as an MLE module call, publishes the JavaScript function helloWorld(). It can then be used just like any other PL/SQL procedure. The following snippet shows how to invoke the function along with the results:

```
SET SERVEROUTPUT ON
BEGIN
    helloWorld_proc;
END;
/
```

Result:

```
Hello World, this is a JS module
```

In addition to custom-built JavaScript modules as shown in the provided code, it is possible to load third-party JavaScript modules into the database. Note that Oracle recommends performing a security screening of third-party code according to industry best practice.

See Also:

- MLE JavaScript Modules and Environments for details about MLE modules and environments
- MLE Security for more information about MLE security features and recommendations



1.3.5 About MLE Execution Contexts

An MLE execution context encapsulates all runtime state associated with the execution of JavaScript code in MLE.

An MLE execution context is a standalone, isolated runtime environment, designed to contain all runtime state associated with the execution of JavaScript code. Runtime state includes global variables as well as the state of the language environment.



An MLE execution context corresponds to an ECMAScript Execution Context for JavaScript.

MLE uses execution contexts in two different scenarios,

- With dynamic MLE execution, where developers may create and use dynamic MLE contexts explicitly
- For calls from SQL and PL/SQL to functions exported by an MLE module

The specific properties of an execution context are configured by an MLE environment. For example, an environment can be used to specify language options and to make MLE modules available for import in an execution context.

Properties of dynamic MLE contexts are determined by the environment used at the moment the execution context is created. Contexts for MLE module calls from SQL or PL/SQL are created implicitly on demand. Here, the properties are determined by the environment referenced in the MLE call specification at the moment of context creation.

MLE modules never share an execution context with other modules or dynamic MLE snippets. Additionally, separate execution contexts are used when code from the same MLE module is executed on behalf of different users. You have explicit control over which execution context is used for each dynamic MLE snippet, with each execution context running code on behalf of a single user.

MLE creates a dedicated execution context for each combination of MLE module and environment. Two call specifications that specify either different modules or different environments are executed in separate module contexts.

There is no limit to how many dynamic MLE execution contexts can be created in a session, or how they are shared across different code snippets. Code snippets in JavaScript share all global variables with other code snippets running in the same execution context.



See Also:

- Specifying Environments for MLE Modules for more information about MLE environments
- Execution Contexts for information about how execution contexts are used to enforce runtime state isolation

1.4 Introduction to Debugging JavaScript Code

MLE allows you to debug your JavaScript code by conveniently and efficiently collecting runtime state during program execution.

After your MLE code has finished executing, debug data collected can be used to analyze program behavior and discover and fix bugs. This form of debugging is known as post-execution debugging.

The post-execution debug option allows you to instrument your code with debugpoints. Debugpoints allow for the logging of program state conditionally or unconditionally, including values of individual variables as well as execution snapshots. Debugpoints are specified as JSON documents separate from the application code. No change to the application code is necessary for debugpoints to fire.

When activated, debug information is collected according to the debug specification and can be fetched for later analysis by a wide range of tools thanks to its standard format.



Post-Execution Debugging of MLE JavaScript Modules for more details about post-execution debugging with MLE



MLE JavaScript Modules and Environments

A JavaScript module is a unit of MLE's language code stored in the database as a schema object.

Storing code within the database is one of the main benefits of using JavaScript in MLE: rather than having to manage a fleet of application servers each with their own copy of the application, the database takes care of this for you.

In addition, Data Guard replication ensures that the same code is present in both production and all physical standby databases. Configuration drift, a common problem bound to occur when invoking the disaster recovery location, can be mitigated against.

A JavaScript module in MLE is equivalent to an ECMAScript 6 module. The terms MLE module and JavaScript module are used interchangeably. The contents are specific to JavaScript and can be managed using Data Definition Language (DDL) commands.

In traditional JavaScript environments, additional information is often passed to the runtime using directives or configuration scripts. In MLE, this can be achieved using MLE environments, an additional metadata structure complementing MLE modules. MLE environments are also used for name resolution of JavaScript module imports. Name resolution is crucial for maintaining code and separating it into various modules to be used with MLE.



Developer.mozilla.org for more information about JavaScript modules

Topics

- Using JavaScript Modules in MLE
 JavaScript modules can be used in several different ways and can be managed using a set of Data Definition Language (DDL) commands.
- Specifying Environments for MLE Modules
 MLE environments are schema objects in the database. Their functionality and
 management methods are described.

2.1 Using JavaScript Modules in MLE

JavaScript modules can be used in several different ways and can be managed using a set of Data Definition Language (DDL) commands.

JavaScript code provided in MLE modules can be used in the following ways:

 JavaScript functions exported by an MLE modules can be published by creating a call specification known as an MLE module call. This allows the function to be called directly from SQL and PL/SQL.

- Functionality exported by a JavaScript MLE module can be imported in other MLE JavaScript modules.
- Code snippets in DBMS_MLE can import modules for dynamic invocation of JavaScript.

Before a user can create and execute MLE modules, several privileges must be granted.

See Also:

- Overview of Importing MLE JavaScript Modules for more information about module calls
- Overview of Dynamic MLE Execution for more information about DBMS MLE and dynamic invocation of JavaScript code in the database
- System and Object Privileges Required for Working with JavaScript in MLE for more information about MLE-specific privileges

Topics

- Managing JavaScript Modules in the Database SQL allows the creation of MLE modules as schema objects, assuming the necessary privileges are in place.
- Preparing JavaScript code for MLE Module Calls
 JavaScript modules in MLE follow the ECMAScript 6 standard for modules.
 Functions and variables expected to be consumed by users of the MLE module must be exported.
- Additional Options for Providing JavaScript Code to MLE
 The JavaScript source code of an MLE module can be specified inline with
 PL/SQL but can also be provided using a BFILE, BLOB, or CLOB, in which case
 the source file must be UTF8 encoded.
- Specifying Module Version Information and Providing JSON Metadata
 MLE modules may carry optional metadata in the form of a version string and freeform JSON-valued metadata.
- Drop JavaScript Modules
 The DROP MLE MODULE DDL statement is used for dropping an MLE module.
- Alter JavaScript Modules
 Attributes of an MLE module can be assigned or altered using the ALTER MLE MODULE statement.
- Overview of Built-in JavaScript Modules
 MLE provides a set of built-in JavaScript modules that are available for import in any execution context.
- Dictionary Views Related to MLE JavaScript Modules
 The Data Dictionary includes details about JavaScript modules.



2.1.1 Managing JavaScript Modules in the Database

SQL allows the creation of MLE modules as schema objects, assuming the necessary privileges are in place.

At a minimum, you need the CREATE MLE MODULE privilege to create or replace an MLE module in your own schema. Additionally, you must have the execute privilege on the target JavaScript language object.

See Also:

- System and Object Privileges Required for Working with JavaScript in MLE for more information about MLE-specific privileges
- Oracle Database Security Guide for more details about privileges and roles in Oracle Database

Topics

- Naming JavaScript Modules
 - Each JavaScript module name must be unique in the schema that it is created in. Unless a fully qualified name is used, the current user's schema is used.
- Creating JavaScript Modules in the Database
 JavaScript modules are created in the database using the CREATE MLE MODULE DDL
 statement, specifying name and source code of the MLE module.
- Storing JavaScript Code in Databases Using Single-Byte Character Sets
 Character set standards and things to remember when using a single-byte character set with MLE.
- Code Analysis

MLE does not perform any checks on JavaScript code for correctness or coding best practice. Thus, developers should use their linting tool of choice to perform analysis before executing the CREATE MLE MODULE command.

2.1.1.1 Naming JavaScript Modules

Each JavaScript module name must be unique in the schema that it is created in. Unless a fully qualified name is used, the current user's schema is used.

As with other schema object identifiers, the module name is case-sensitive if enclosed in double quotation marks. If the enclosing quotation marks are omitted, the name is implicitly converted to uppercase.

When choosing a unique name, note that MLE objects share the namespace with tables, views, materialized views, sequences, private synonyms, PL/SQL packages, functions, procedures, and cache groups.



2.1.1.2 Creating JavaScript Modules in the Database

JavaScript modules are created in the database using the CREATE MLE MODULE DDL statement, specifying name and source code of the MLE module.

As soon as an MLE module has been created, it is persisted in the database dictionary. This is one of the differences when compared to dynamic execution of JavaScript code using DBMS MLE.

CREATE MLE MODULE (without the OR REPLACE clause) throws an error if an MLE module with the given name already exists. With CREATE OR REPLACE MLE MODULE, the existing module is replaced if it exists, otherwise a new one is created. When an MLE module is replaced, privileges to it do not need to be re-granted.

For those who are familiar with PL/SQL, note that this is exactly the same behavior experienced with PL/SQL program units.

If you do not wish to replace an existing module in the event the module name is already in use, you can use the IF NOT EXISTS clause rather than CREATE OR REPLACE. The syntax for this variation is shown in Example 2-1. The IF NOT EXISTS and OR REPLACE clauses are mutually exclusive.



Oracle Database Development Guide for more information about using the IF [NOT] EXISTS syntax

Example 2-1 Creating a JavaScript Module in the Database

This example demonstrates the creation of an MLE module and the export of a simple JavaScript function.

CREATE MLE MODULE IF NOT EXISTS po module LANGUAGE JAVASCRIPT AS



The first line of this code block specifies the JavaScript module name as po_module. The remaining lines define the actual JavaScript code. Note that in line with the ECMAScript standard, the export keyword indicates the function to be exported to potential callers of the module. MLE accepts code adhering to the ECMAScript 2022 standard.

2.1.1.3 Storing JavaScript Code in Databases Using Single-Byte Character Sets

Character set standards and things to remember when using a single-byte character set with MLE.

JavaScript is encoded in Unicode. The Unicode Standard is a character encoding system that defines every character in most of the spoken languages in the world. It was developed to overcome limitations of other character-set encodings.

Oracle recommends creating databases using the AL32UTF8 character set. Using the AL32UTF8 character set in the database ensures the use of the latest version of the Unicode Standards and minimizes the potential for character-set conversion errors.

In case your database still uses a single-byte character set such as US7ASCII, WE8ISO8859-n, or WE8MSWIN1252, you must be careful not to use Unicode features in MLE JavaScript code. This is no different than handling other types of input data with such a database.



Oracle Database Globalization Support Guide for more details about the Unicode Standard

2.1.1.4 Code Analysis

MLE does not perform any checks on JavaScript code for correctness or coding best practice. Thus, developers should use their linting tool of choice to perform analysis before executing the CREATE MLE MODULE command.

When creating MLE modules in the database, you should use a well-established toolchain in the same way other JavaScript projects are governed. In this sense, the call to CREATE MLE MODULE can be considered a deployment step, similar to deploying a server application. Code checking should be performed during a build step, for example by a continuous integration and continuous deployment (CI/CD) pipeline, prior to deployment.

It is considered an industry best practice to process code with a tool called a *linter* before checking it into a source-code repository. As with any other development project, you are free to choose the best option for yourself and your team. Some potential options include ESLint, JSHint, JSLint, and others that perform static code analysis to flag syntax errors, bugs, or otherwise problematic code. They can also be used to enforce a certain coding style. Many integrated development environments (IDEs) provide linting as a built-in feature, invoking the tool as soon as a file is saved to disk and flagging any issues.

In addition to executing linting dynamically, it is possible to automate the code analysis using highly automated DevOps environments to invoke linting as part of a build pipeline. This step usually occurs prior to submitting the JavaScript module to the database.



The aim is to trap as many potential issues as possible before they can produce problems at runtime. Unit tests can help further mitigate these risks and their inclusion into the development process have become an industry best practice. Regardless of the method you choose, the code analysis step occurs prior to submitting the JavaScript module to the database.

2.1.2 Preparing JavaScript code for MLE Module Calls

JavaScript modules in MLE follow the ECMAScript 6 standard for modules. Functions and variables expected to be consumed by users of the MLE module must be exported.

Those variables and functions *not* exported are considered private in the module. Example 2-3 demonstrates the use of both public and private functions in an MLE JavaScript module.

An ECMAScript module can import other ECMAScript modules using import statements or dynamic import calls. This functionality is present in MLE as well. Complementary metadata to MLE modules is provided in MLE environments.

Note that console output in MLE is facilitated using the console object. By default, anything written to console.log() is routed to <code>DBMS_OUTPUT</code> and will end up on the screen.

JavaScript code like that in Example 2-1 cannot be accessed from SQL or PL/SQL without the help of call specifications. For now, you can think of a call specification as a PL/SQL program unit (function, procedure, or package) where its PL/SQL body is replaced with a reference to the JavaScript module and function, as shown in Example 2-2. For more information about call specifications, see MLE JavaScript Functions.



Using MLE Environments for Import Resolution

Example 2-2 Create a Call Specification for a Public Function

This example uses the module po_module created in Example 2-1. A call specification for orderValue(), the only function exported in po module, can be written as follows:

```
CREATE OR REPLACE FUNCTION order_value(
    p_line_items JSON
) RETURN NUMBER AS
MLE MODULE po_module
SIGNATURE 'orderValue';
/
```

Once the function is created, it is possible to calculate the value of a given purchase order:

SELECT



```
po.po_document.PONumber,
  order_value(po.po_document.LineItems[*]) order_value
FROM
  j_purchaseorder po;
```

Result:

PONUMBER	ORDER	_VALUE
1600		279.3
672		359.5

Example 2-3 Public and Private Functions in a JavaScript Module

In addition to public (exported) functions, it is possible to add functions private to the module. In this example, the calculation of the value is taken out of the map() function and moved to a separate function (refactoring).

The first function in the following code, <code>lineItemValue()</code>, is considered private, whereas the second function, <code>orderValue()</code>, is public. The export keyword is provided at the end of this code listing but can also appear as a prefix for variables and functions, as seen in <code>Example 2-1</code>. Both variations are valid JavaScript syntax.

```
CREATE OR REPLACE MLE MODULE po module LANGUAGE JAVASCRIPT AS
/**
* calculate the value of a given line item. Factored out of the public
* function to allow for currency conversions in a later step
* @param {number} unitPrice - the price of a single article
* @param {number} quantity - the quantity of articles ordered
* @returns {number} the monetary value of the line item
function lineItemValue(unitPrice, quantity) {
    return unitPrice * quantity;
}
* get the value of all line items in an order
* @param {array} lineItems - all the line items in a purchase order
* @returns {number} the total value of all line items in a purchase order
function orderValue(lineItems) {
    return lineItems
                .map(x => lineItemValue(x.Part.UnitPrice, x.Quantity))
                .reduce(
                    (accumulator, currentValue) => accumulator +
currentValue, 0
                );
}
```



```
export { orderValue }
/
```

2.1.3 Additional Options for Providing JavaScript Code to MLE

The JavaScript source code of an MLE module can be specified inline with PL/SQL but can also be provided using a BFILE, BLOB, or CLOB, in which case the source file must be UTF8 encoded.

Creating MLE modules using the BFILE clause can cause problems with logical replication such as GoldenGate. In order for the DDL command to succeed on the target database, the same directory must exist on the target database. Furthermore, the same JavaScript file must be present in this directory. Failure to adhere to these conditions will cause the call to create the MLE module on the target database to fail.

Example 2-4 Providing JavaScript Source Code Using a BFILE

In this example, <code>JS_SRC_DIR</code> is a database directory object mapping to a location on the local file system containing the module's source code in a file called <code>myJavaScriptModule.js</code>. When loading the file from the directory location, MLE stores the source code in the dictionary. Subsequent calls to the MLE module will not cause the source code to be refreshed from the disk. If there is a new version of the module stored in <code>myJavaScriptModule.js</code>, it must be deployed using another call to <code>CREATE OR REPLACE MLE MODULE</code>.

```
CREATE MLE MODULE mod_from_bfile
LANGUAGE JAVASCRIPT
USING BFILE(JS_SRC_DIR,'myJavaScriptModule.js');
/
```

2.1.4 Specifying Module Version Information and Providing JSON Metadata

MLE modules may carry optional metadata in the form of a version string and freeform JSON-valued metadata.

Both kinds of metadata are purely informational and do not influence the behavior of MLE. They are stored alongside the module in the data dictionary.

The VERSION flag can be used as an internal reminder about what version of the code is deployed. The information stored in the VERSION field allows developers and administrators to identify the code in the version control system.

The format for JSON metadata is not bound to a schema; anything useful or informative can be added by the developer. In case the MLE module is an aggregation of sources created by a tool such as rollup.js or webpack, it can be useful to store the associated package-lock.json file alongside the module.

The metadata field can be used to create a software bill of material (SBOM), allowing security teams and administrators to track information about deployed packages, especially in the case where third-party modules are used.



Tracking dependencies and vulnerabilities in the upstream repository supports easier identification of components in need of update after security vulnerabilities have been reported.



- Dictionary Views Related to MLE JavaScript Modules
- Software Bill of Material for more information about using the metadata field to store a SBOM

Example 2-5 Specification of a VERSION string in CREATE MLE MODULE

```
CREATE OR REPLACE MLE MODULE version_mod
  LANGUAGE JAVASCRIPT
  VERSION '1.0.0.1.0'
AS
  export function sq(num) {
   return num * num;
}
//
```

Example 2-6 Addition of JSON Metadata to the MLE Module

This example uses the module version mod, created in Example 2-5.

```
ALTER MLE MODULE version_mod
SET METADATA USING CLOB

(SELECT

'{
    "name": "devel",
    "lockfileVersion": 2,
    "requires": true,
    "packages": {}
}'
)
/
```

2.1.5 Drop JavaScript Modules

The DROP MLE MODULE DDL statement is used for dropping an MLE module.

The $\[DROP\]$ statement specifies the name, and optionally the schema of the module to be dropped. If a schema is not specified, the schema of the current user is assumed.

Attempting to drop an MLE module that does not exist causes an error to be thrown. In cases where this is not desirable, the IF EXISTS clause can be used. The DROP MLE MODULE command is silently skipped if the indicated MLE module does not exist.

Example 2-7 Drop an MLE Module

DROP MLE MODULE unused mod;

Example 2-8 Drop an MLE Module Using IF EXISTS

DROP MLE MODULE IF EXISTS unused mod;

2.1.6 Alter JavaScript Modules

Attributes of an MLE module can be assigned or altered using the ALTER MLE MODULE statement.

The ALTER MLE MODULE statement specifies the name, and optionally the schema of the module to be altered. If the module name is not prefixed with a schema, the schema of the current user is assumed.

Example 2-9 Alter an MLE Module

```
ALTER MLE MODULE change_mod
SET METADATA USING CLOB(SELECT'{...}');
```

2.1.7 Overview of Built-in JavaScript Modules

MLE provides a set of built-in JavaScript modules that are available for import in any execution context.

Built-in modules are not deployed to the database as user-defined MLE modules, but are included as part of the MLE runtime. In particular, MLE provides the following three built-in JavaScript modules:

- mle-js-oracledb is the JavaScript MLE SQL Driver.
- mle-js-bindings provides functionality to import and export values from the PL/SQL engine.
- mle-js-plsqltypes provides definitions for the PL/SQL wrapper types. For example, JavaScript types that wrap PL/SQL and SQL types like OracleNumber.
- mle-js-fetch provides a partial Fetch API polyfill, allowing developers to invoke external resources.
- mle-encode-base64 contains code to work with base64-encoded data.

These modules can be used to interact with the database and provide type conversions between the JavaScript engine and database engine.



Server-Side JavaScript API Documentation for more information about the built-in JavaScript modules

2.1.8 Dictionary Views Related to MLE JavaScript Modules

The Data Dictionary includes details about JavaScript modules.

Topics

USER_SOURCE

Each JavaScript module's source code is externalized using the [USER | ALL | DBA | CDB] SOURCE dictionary views.

USER MLE MODULES

Metadata pertaining to JavaScript MLE modules are found in [USER | ALL | DBA | CDB] MLE MODULES.

2.1.8.1 USER SOURCE

Each JavaScript module's source code is externalized using the [USER | ALL | DBA | CDB] SOURCE dictionary views.

Modules created with references to the file system using the BFILE operator show the code at the time of the module's creation.

For more information about * SOURCE, see Oracle Database Reference.

Example 2-10 Externalize JavaScript Module Source Code

```
SELECT
    line,
    text
FROM
    USER_SOURCE
WHERE
    name = 'PO_MODULE';
```

Example output:

```
LINE TEXT

----

1 /**

2 * calculate the value of a given line item. Factored out of the public

3 * function to allow for currency conversions in a later step

4 * @param {number} unitPrice - the price of a single article

5 * @param {number} quantity - the quantity of articles ordered

6 * @returns {number} the monetary value of the line item

7 */

8 function lineItemValue(unitPrice, quantity) {

9 return unitPrice * quantity;

10 }

11
```

```
12
        13 /**
        14 * get the value of all line items in an order
       15 * @param {array} lineItems - all the line items in a
purchase order
       16 * @returns {number} the total value of all line items in a
purchase order
       17 */
        18 export function orderValue(lineItems) {
        20
             return lineItems
        21
                          .map( x => lineItemValue(x.Part.UnitPrice,
x.Quantity) )
                          .reduce(
        22
        2.3
                              (accumulator, currentValue) =>
accumulator + currentValue, 0
        24
        25 }
```

2.1.8.2 USER MLE MODULES

Metadata pertaining to JavaScript MLE modules are found in [USER | ALL | DBA | CDB]_MLE_MODULES.

Any JSON metadata specified, version information, as well as language, name, and owner can be found in this view.

For more information about *_MLE_MODULES, see *Oracle Database Reference*.

Example 2-11 Find MLE Modules Defined in a Schema

```
SELECT MODULE_NAME, VERSION, METADATA FROM USER_MLE_MODULES
WHERE LANGUAGE_NAME='JAVASCRIPT'
/
```

Example output:

MODULE_NAME	VERSION	METADATA
MY_MOD01	1.0.0.1	
MY_MOD02	1.0.1.1	
MY MOD03		

2.2 Specifying Environments for MLE Modules

MLE environments are schema objects in the database. Their functionality and management methods are described.

MLE environments complement MLE modules and allow you to do the following:

- Set language options to customize the JavaScript runtime in its execution context
- Enable specific MLE modules to be imported

Manage name resolution and the import chain

Topics

Creating MLE Environments in the Database
 The SQL DDL supports the creation of MLE environments.

Dropping MLE Environments

MLE environments that are no longer needed can be dropping using the DROP MLE ENV command.

Modifying MLE Environments

Existing MLE environments can be modified using the ALTER MLE ENV command.

Dictionary Views Related to MLE JavaScript Environments
 Details about MLE environments are available in these families of views: USER_MLE_ENVS and USER MLE ENV IMPORTS.

2.2.1 Creating MLE Environments in the Database

The SQL DDL supports the creation of MLE environments.

Just like MLE modules, MLE environments are schema objects in the database, persisted in the data dictionary.

At a minimum, you must have the CREATE MLE MODULE privilege to create or replace an MLE environment in your own schema.

See Also:

- System and Object Privileges Required for Working with JavaScript in MLE for more information about the privileges necessary to create and execute JavaScript code in MLE
- Oracle Database Security Guide for details about privileges and roles in Oracle Database

Topics

Naming MLE Environments

Each JavaScript environment's name must be unique in the schema it is created in. Unless a fully qualified name is used, the current user's schema is used.

Creating an Empty MLE Environment

The DDL statement CREATE MLE ENV can be used to create an MLE environment.

- Creating an Environment as a Clone of an Existing Environment
 If needed, a new environment can be created as a point-in-time copy of an existing environment.
- Using MLE Environments for Import Resolution
 It is possible to import functionality exported by one JavaScript module into another using the import statement.
- Providing Language Options
 MLE allows the customization of JavaScript's runtime by setting language-specific options in MLE environments.

2.2.1.1 Naming MLE Environments

Each JavaScript environment's name must be unique in the schema it is created in. Unless a fully qualified name is used, the current user's schema is used.

As with other schema object identifiers, the name is case-sensitive if enclosed in double quotation marks. If the enclosing quotation marks are omitted, the name is implicitly converted to uppercase.

MLE environments cannot contain import mappings that conflict with the names of the MLE built-in modules (mle-js-oracledb, mle-js-bindings, mle-js-plsqltypes, mle-js-fetch, and mle-encode-base64). If you attempt to add such a mapping using either the CREATE MLE ENV OF ALTER MLE ENV DDL, the operation fails with an error.

2.2.1.2 Creating an Empty MLE Environment

The DDL statement CREATE MLE ENV can be used to create an MLE environment.

In its most basic form, an environment can be created empty as shown in the following snippet:

CREATE MLE ENV MyEnv

Subsequent calls to ALTER MLE ENV can be used to add properties to the environment.

Just like with MLE modules, it is possible to append the OR REPLACE clause to instruct the database to replace an existing MLE environment rather than throwing an error.

Furthermore, the IF NOT EXISTS clause can be used instead of the OR REPLACE clause to prevent the creation of a new MLE environment in the case that one already exists with the same name. In this case, the statement used to create the environment changes to the following:

CREATE MLE ENV IF NOT EXISTS MyEnv



The IF NOT EXISTS and OR REPLACE clauses are mutually exclusive.



Modifying MLE Environments for information about editing existing environments



2.2.1.3 Creating an Environment as a Clone of an Existing Environment

If needed, a new environment can be created as a point-in-time copy of an existing environment.

The new environment inherits all settings from its source. Subsequent changes to the source are not propagated to the clone. A clone can be created as shown in the following statement:

CREATE MLE ENV MyEnvDuplicate CLONE MyEnv

2.2.1.4 Using MLE Environments for Import Resolution

It is possible to import functionality exported by one JavaScript module into another using the import statement.

The separation of code allows for finer control over changes and the ability to write more reusable code. Simplified code maintenance is another positive effect of this approach.

Only those identifiers marked with the export keyword are eligible for importing.

Modules attempting to import functionality from other modules stored in the database require MLE environments in order to perform name resolution. To create an MLE environment with that information, the IMPORTS clause must be used. Example 2-12 demonstrates how a mapping is created between the identifier po_module and JavaScript module PO_MODULE, created in Example 2-1.

Multiple imports can be provided as a comma-separated list. In Example 2-12, the first parameter in single quotation marks is known as the import name. The import name is used by another module's import statement. In this case, 'po_module' is the import name and refers to the module of the same name.



The import name does not have to match the module name. Any valid JavaScript identifier can be used. The closer the import name matches the module name it refers to, the easier it is to identify the link between the two.

The CREATE MLE ENV command fails if a module referenced in the IMPORTS clause does not exist or is not accessible to you.

Built-in JavaScript modules can be imported directly without having to specify additional MLE environments.



Overview of Built-in JavaScript Modules for more information about built-in modules



Example 2-12 Map Identifier to JavaScript Module

```
CREATE OR REPLACE MLE ENV
    po_env
IMPORTS (
    'po_module' MODULE PO_MODULE
);
```

Example 2-13 Import Module Functionality

```
CREATE OR REPLACE MLE MODULE import_example_module LANGUAGE JAVASCRIPT AS
```

```
import * as po from "po module";
* use po module's getValue() function to calculate the value of
^{\star} a purchase order. In later chapters, when discussing the MLE
* JavaScript SQL driver the hard-coded value used as the PO will
* be replaced by calls to the database
* @returns {number} the value of all line items in the purchase order
*/
export function purchaseOrderValue() {
    const purchaseOrder = {
        "PONumber": 1600,
        "Reference": "ABULL-20140421",
        "Requestor": "Alexis Bull",
        "User": "ABULL",
        "CostCenter": "A50",
        "ShippingInstructions": {
            "name": "Alexis Bull",
            "Address": {
                "street": "200 Sporting Green",
                "city": "South San Francisco",
                "state": "CA",
                "zipCode": 99236,
                "country": "United States of America"
            },
            "Phone": [
                {
                    "type": "Office",
                    "number": "909-555-7307"
                },
                    "type": "Mobile",
                    "number": "415-555-1234"
            ]
        },
        "Special Instructions": null,
        "AllowPartialShipment": true,
        "LineItems": [
```



```
{
                "ItemNumber": 1,
                "Part": {
                    "Description": "One Magic Christmas",
                    "UnitPrice": 19.95,
                    "UPCCode": 13131092899
                },
                "Quantity": 9.0
            },
                "ItemNumber": 2,
                "Part": {
                    "Description": "Lethal Weapon",
                    "UnitPrice": 19.95,
                    "UPCCode": 85391628927
                "Quantity": 5.0
            }
        1
    };
    return po.orderValue(purchaseOrder.LineItems);
}
purchaseOrderValue
CREATE FUNCTION purchase_order_value
RETURN NUMBER AS
MLE MODULE import_example_module
ENV po env
SIGNATURE 'purchaseOrderValue';
SELECT purchase_order_value;
Result:
PURCHASE ORDER VALUE
______
               279.3
```

2.2.1.5 Providing Language Options

MLE allows the customization of JavaScript's runtime by setting language-specific options in MLE environments.

Any options specified in the MLE environment take precedence over the default settings.

Multiple language options can be provided as a comma-separated list of '<key>=<value>' strings. The following snippet demonstrates how to enforce JavaScript's strict mode.

```
CREATE MLE ENV MyEnvOpt
LANGUAGE OPTIONS 'js.strict=true';
```

Changes made to the language options of an environment are not propagated to execution contexts that have already been created using the environment. For changes to take effect for existing contexts, the contexts need to be dropped and recreated.



White space characters are not allowed between the key, equal sign, and value.

Topics

JavaScript Language Options
 A full list of JavaScript language options available to be used with MLE are included.

2.2.1.5.1 JavaScript Language Options

A full list of JavaScript language options available to be used with MLE are included.

Table 2-1 JavaScript Language Options

Language Option	Accepted Value Type	Default	Description
js.strict	boolean	false	Enforce strict mode.
js.console	boolean	true	Provide console global property.
js.polyglot-builtin	boolean	true	Provide Polyglot global property.

2.2.2 Dropping MLE Environments

MLE environments that are no longer needed can be dropping using the ${\tt DROP}\ {\tt MLE}\ {\tt ENV}$ command.

The following snippet demonstrates a basic example of dropping an MLE module:

```
DROP MLE ENV myOldEnv;
```

As with MLE modules, the IF EXISTS clause prevents an error if the named MLE environment does not exist, as shown in the following snippet:

DROP MLE ENV IF EXISTS myOldEnv;



2.2.3 Modifying MLE Environments

Existing MLE environments can be modified using the ALTER MLE ENV command.

It is possible to modify language options and the imports clause.

Topics

- Altering Language Options
 You can modify language options provided to an MLE module.
- Modifying Module Imports
 In the context of MLE module imports, the ALTER MLE ENV command allows you to add additional imports as well as modify and drop existing imports.

2.2.3.1 Altering Language Options

You can modify language options provided to an MLE module.

Use the ALTER MLE ENV clause to modify language options, as shown in the following snippet:

```
ALTER MLE ENV MyEnvOpt

SET LANGUAGE OPTIONS 'js.strict=false';
```

2.2.3.2 Modifying Module Imports

In the context of MLE module imports, the ALTER MLE ENV command allows you to add additional imports as well as modify and drop existing imports.

Imports not specified during an environment's creation can be added to existing MLE environments using the ADD IMPORTS clause. Import names, once defined, are static and must be dropped before they can be added as desired. Assuming you have run a new CREATE MLE DDL to replace IMPORT_EXAMPLE_MODULE from Example 2-1 with the module name IMPORT_EXAMPLE_MODULE_V2, the following statement will run successfully:

```
ALTER MLE ENV po_env
ADD IMPORTS (
    'import_example' MODULE IMPORT_EXAMPLE_MODULE_V2
);
```

Any imports no longer needed can be dropped using the DROP IMPORTS clause:

```
ALTER MLE ENV po_env DROP IMPORTS('import_example');
```

The case of the import identifier must match that in the data dictionary's USER MLE ENV IMPORTS view.



2.2.4 Dictionary Views Related to MLE JavaScript Environments

Details about MLE environments are available in these families of views: USER MLE ENVS and USER MLE ENV IMPORTS.

In addition to the USER prefix, these views exist in all namespaces: CDB, DBA, ALL, and USER.

Topics

USER MLE ENVS

The USER_MLE_ENVS view lists all MLE environments available to you along with the defined language options.

USER MLE ENV IMPORTS

The [USER | ALL | DBA | CDB]_MLE_ENV_IMPORTS family of views lists imported modules.

2.2.4.1 USER MLE ENVS

The $\tt USER_MLE_ENVS$ view lists all MLE environments available to you along with the defined language options.

For more information about $*_{\texttt{MLE}_\texttt{ENVS}}$, see *Oracle Database Reference*.

Example 2-14 List Available MLE Environments Using USER_MLE_ENVS

```
SELECT ENV_NAME, LANGUAGE_OPTIONS
FROM USER_MLE_ENVS
WHERE ENV_NAME='MYENVOPT'
/
```

Example SQL*Plus output:

ENV_OWNER	ENV_NAME	LANGUAGE_OPTIONS
JSDEV01	MYENVOPT	js.strict=true

2.2.4.2 USER_MLE_ENV_IMPORTS

The [USER | ALL | DBA | CDB]_MLE_ENV_IMPORTS family of views lists imported modules.

MLE environments are the key enablers for resolving names of imported modules. Example 2-15 demonstrates a query against <code>USER_MLE_ENV_IMPORTS</code> to list <code>IMPORT NAME</code>, <code>MODULE OWNER</code>, and <code>MODULE NAME</code>.

For more information about * MLE ENV IMPORTS, see Oracle Database Reference

Example 2-15 List Module Import Information Using USER_MLE_ENV_IMPORTS

```
SELECT IMPORT_NAME, MODULE_OWNER, MODULE_NAME FROM USER MLE ENV IMPORTS
```



```
WHERE ENV_NAME='MYFACTORIALENV';
```

SQL*Plus output:

IMPORT_NAME	MODULE_OWNER	MODULE_NAME
FACTORIAL_MOD	DEVELOPER1	FACTORIAL_MOD



Overview of Dynamic MLE Execution

Dynamic MLE execution allows developers to invoke JavaScript snippets via the DBMS_MLE package without storing the JavaScript code in the database.

As an alternative to executing JavaScript code using modules, MLE provides the option of dynamic execution. With dynamic execution, no JavaScript code is stored in the data dictionary. Instead, you can invoke snippets of JavaScript code via the <code>DBMS_MLE</code> package.

See Also:

- Server-Side JavaScript API Documentation for information about built-in module mle-js-bindings, used to exchange values between PL/SQL and JavaScript
- Oracle Database PL/SQL Packages and Types Reference for more information about the DBMS MLE package

Topics

- About Dynamic JavaScript Execution
 - Developers can run JavaScript dynamically either inline or by loading files via DBMS_MLE. Dynamic MLE execution provides an additional method for using JavaScript to interact with the Oracle Database, as an alternative to using MLE modules.
- Dynamic Execution Workflow
 The steps required to perform dynamic MLE execution are described.
- Returning the Result of the Last Execution
 Use the result argument to get the outcome of the last execution.

3.1 About Dynamic JavaScript Execution

Developers can run JavaScript dynamically either inline or by loading files via <code>DBMS_MLE</code>. Dynamic MLE execution provides an additional method for using JavaScript to interact with the Oracle Database, as an alternative to using MLE modules.

The DBMS_MLE package allows users to execute JavaScript code inside the Oracle Database and seamlessly exchange data between PL/SQL and JavaScript. The JavaScript code itself can execute PL/SQL through built-in JavaScript modules. JavaScript data types are automatically mapped to Oracle Database data types and vice versa.

Developers can provide JavaScript code either as the value of a VARCHAR2 variable or, in case of larger amounts of code, as a Character Large Object (CLOB). The JavaScript code is passed to the DBMS MLE package where it is evaluated and executed.

Considering that <code>DBMS_MLE</code> is a PL/SQL package, there is mix of JavaScript and PL/SQL when dynamically executing code using <code>DBMS_MLE</code>, for example, in the following cases:

- Setup tasks such as providing the JavaScript code require an interaction with the PL/SQL layer.
- JavaScript code is executed by calling a function in DBMS MLE.
- After JavaScript code completes execution, any errors that have been encountered are passed back to PL/SQL.

3.2 Dynamic Execution Workflow

The steps required to perform dynamic MLE execution are described.

Before a user can create and execute JavaScript code using <code>DBMS_MLE</code>, several privileges must be granted. For information about required privileges, see System and Object Privileges Required for Working with JavaScript in MLE.

The execution workflow for JavaScript code using DBMS MLE is as follows:

- Create an execution context
- 2. Provide JavaScript code either using a VARCHAR2 or CLOB variable
- Execute the code, optionally passing variables between the PL/SQL and MLE engines
- Close the execution context

As with any code, it is considered an industry best practice to deal with unexpected conditions. You can do this in the JavaScript code itself using standard JavaScript exception handling features or in PL/SQL.

Topics

- Providing JavaScript Code Inline
 Suggested method for providing JavaScript code inline by embedding it within a PL/SQL block.
- Loading JavaScript Code from Files
 The method for using a BFILE operator to read in a CLOB is described.

3.2.1 Providing JavaScript Code Inline

Suggested method for providing JavaScript code inline by embedding it within a PL/SQL block.

One option to load JavaScript code is to use a quoting operator, commonly referred to as the q-quote operator. The use of this alternative quoting operator is suggested as the preferred method to provide JavaScript code inline with PL/SQL code whenever possible.

Example 3-1 Using the Q-Quote Operator to Provide JavaScript Code Inline with PL/SOL

```
DECLARE
    l_ctx         dbms_mle.context_handle_t;
    l_snippet CLOB;
BEGIN
    l_ctx := dbms_mle.create_context();
    l snippet := q'~
```



```
// the q-quote operator allows for much more readable code
console.log(`The use of the q-quote operator`);
console.log(`greatly simplifies provision of code inline`);
~';
   dbms_mle.eval(l_ctx, 'JAVASCRIPT', l_snippet);
   dbms_mle.drop_context(l_ctx);
EXCEPTION
   WHEN OTHERS THEN
        dbms_mle.drop_context(l_ctx);
   RAISE;
END;
/
Result:
The use of the q-quote operator
greatly simplifies provision of code inline
```

3.2.2 Loading JavaScript Code from Files

The method for using a BFILE operator to read in a CLOB is described.

If you plan to use a linter to conduct code analysis, providing JavaScript code in line with PL/SQL may not be your best option for dynamic execution. Another method for providing JavaScript code is to read a CLOB by means of a BFILE operator. This way PL/SQL and JavaScript code can be cleanly separated.



Oracle Database SecureFiles and Large Objects Developer's Guide for information about Large Objects

Example 3-2 Loading JavaScript code from a BFILE with DBMS_LOB.LOADCLOBFROMFILE()

This example illustrates the use of a BFILE and DBMS LOB.LOADCLOBFROMFILE().

The example assumes that you have read access to a directory named <code>SRC_CODE_DIR</code>. The source code file <code>hello source.js</code> resides in that directory. Its contents are as follows:



```
l warn
                 INTEGER := 0;
BEGIN
    1 ctx := dbms mle.create context();
    dbms lob.createtemporary(lob loc => 1 js, cache => false);
    1_srcode_file := bfilename('SRC CODE DIR', 'hello source.js');
    IF ( dbms lob.fileexists(file loc => 1 srcode file) = 1 ) THEN
        dbms lob.fileopen(file loc => 1 srcode file);
        dbms lob.loadclobfromfile(
                          => 1 js,
                dest lob
                src_bfile => l_srcode_file,
               amount => dbms_lob.getlength(l_srcode_file),
               dest offset => l dest offset,
               src offset => l src offset,
               bfile csid => l csid,
               lang context => 1 lang context,
               warning => l warn
        );
        IF 1 warn = dbms lob.warn inconvertible char THEN
           raise application error(
               -20001,
                'the input file contained inconvertible characters'
           );
        END IF;
        dbms lob.fileclose(l srcode file);
        dbms mle.eval(
           context handle => 1 ctx,
           language_id => 'JAVASCRIPT',
                         => 1 js
           source
        );
        dbms mle.drop context(1 ctx);
    ELSE
       raise application error(
           -20001,
           'The input file does not exist'
        );
   END IF;
EXCEPTION
    WHEN OTHERS THEN
       dbms_mle.drop_context(l_ctx);
       RAISE;
END;
Result:
hello from hello source
```

In some cases, you may need to mix dynamic MLE execution as shown in with MLE modules persisted in the database, as shown in Example 3-3.

Example 3-3 Loading JavaScript Code from a BFILE by Referencing an MLE Module from DBMS_MLE

The code for the JavaScript module is again stored in a file, as seen in Example 3-2. The example assumes that you have read access to a directory named <code>SRC_CODE_DIR</code> and the file name is <code>greeting source.js</code>:

```
export function greeting() {
    return 'hello from greeting_source';
}
```

This example begins by creating an MLE module from BFILE using the contents of the preceding file. Before the module can be used by DBMS_MLE, an environment must be created first, allowing the dynamic portion of the JavaScript code to reference the module.

Dynamic MLE execution does not allow the use of the ECMAScript import keyword. MLE modules must instead be dynamically imported using the async/await interface shown in this example.

```
CREATE OR REPLACE MLE MODULE greet mod
LANGUAGE JAVASCRIPT
USING BFILE(SRC CODE DIR, 'greeting source.js');
CREATE OR REPLACE MLE ENV greet mod env
imports ('greet mod' module greet mod);
DECLARE
    1 ctx dbms mle.context handle t;
    1 snippet CLOB;
BEGIN
    1 ctx := dbms mle.create context(
        environment => 'GREET MOD ENV'
    l snippet := q'~
(async () => {
    let { greeting } = await import('greet mod');
    const message = greeting();
    console.log(message);
})();
~';
    dbms mle.eval(
                 1 ctx,
                 'JAVASCRIPT',
                 1 snippet
    );
    dbms mle.drop context(l ctx);
EXCEPTION
    WHEN OTHERS THEN
        dbms mle.drop context(l ctx);
        RAISE;
```



```
END;
```

Result:

hello from greeting source



Additional Options for Providing JavaScript Code to MLE for information about using BFILES with MLE modules to load JavaScript code

3.3 Returning the Result of the Last Execution

Use the result argument to get the outcome of the last execution.

A variant of the <code>DBMS_MLE.eval()</code> procedure takes an additional CLOB argument, <code>result</code>. Such a call to <code>DBMS_MLE.eval()</code> appends the outcome of the execution of the last statement in the provided dynamic MLE snippet to the CLOB provided as the <code>result</code> parameter.

This option is useful in the implementation of an interactive application, such as a Read-Eval-Print-Loop (REPL) server, to mimic the behavior of a similar REPL session in Node.js.

Example 3-4 Returning the Result of the Last Execution

```
DECLARE
    1 ctx    dbms mle.context handle t;
    1 snippet CLOB;
    l result CLOB;
BEGIN
    dbms lob.createtemporary(
                lob_loc => l_result,
                cache => false,
                dur => dbms lob.session
    );
    l ctx := dbms mle.create context();
    l snippet := q'~
let i = 21;
i *= 2;
    dbms mle.eval(
                context handle => 1 ctx,
                language id => 'JAVASCRIPT',
                source => l_snippet,
result => l_result
    );
```



```
dbms_output.put_line('result: ' || l_result);
   dbms_mle.drop_context(l_ctx);

EXCEPTION
   WHEN OTHERS THEN
        dbms_mle.drop_context(l_ctx);
        RAISE;

END;
/
Result:
```



4

Overview of Importing MLE JavaScript Modules

One of the key benefits of MLE is the ability to store modules of JavaScript code in the Oracle Database. The availability of modules supports self-contained and reusable code, key to developing successful software projects.

MLE modules interact with each other through imports and exports. Specifically, if a module wants to use a functionality provided by another module, the source module must be exported and then imported into the calling module's scope.

Due to a difference in architecture, module imports behave slightly differently in the Oracle Database when compared to other development environments. For example, JavaScript source code used with Node.js is stored in a specific directory structure on disk. Alternatively, MLE modules are stored together with the database, rather than in a file system, so must be referenced in a different manner.

There are two options available for module imports in MLE:

- Importing module functionality into another module
- Importing module functionality into a code snippet to be executed via dynamic MLE (using the DBMS_MLE PL/SQL package)



MLE supports a pure JavaScript implementation. Module exports and imports are facilitated as ECMAScript modules using the export and import keywords. Other JavaScript modularization technologies such as CommonJS and Asynchronous Module Definition (AMD) are not available.



MLE JavaScript Modules and Environments for more information on MLE modules

Topics

- JavaScript Module Hierarchies
 - The use of import names as opposed to file-system location to resolve to MLE modules is described.
- Export Functionality

Commonly exported identifiers in native JavaScript modules include variables, constants, functions, and classes.



Import Functionality

The import keyword allows developers to import functionality that has been exported by a source module.

4.1 JavaScript Module Hierarchies

The use of import names as opposed to file-system location to resolve to MLE modules is described.

A typical Node.js or browser-based workflow requires a module import to be followed by its location in the file system. For example, the following line is a valid module import statement in Node.js:

```
import * as myMath from './myMath.mjs'
```

Used with Node.js, this statement would import myMath's contents into the current scope.

However, because MLE JavaScript modules are stored in the database, there are no file-system paths to be used for identification. Rather, MLE uses import names instead that resolve to the desired module.



As soon as a module import is detected by the JavaScript runtime engine, strict mode is enforced.

Topics

Resolving Import Names Using MLE Environments
 Rather than file-system locations, MLE uses so-called import names instead.

 Import names must be valid JavaScript identifiers, but otherwise can be chosen freely.

4.1.1 Resolving Import Names Using MLE Environments

Rather than file-system locations, MLE uses so-called import names instead. Import names must be valid JavaScript identifiers, but otherwise can be chosen freely.

Example 4-1 Use an MLE Environment to Map an Import Name to a Module

This example shows how you might use an import name for code referencing functionality in module named exports module, which is defined in Example 4-2.

MLE in Oracle Database requires a link between the import name, namedExports, and the corresponding MLE module, named_exports_module, at runtime. Just as with import names, you can choose any valid name for the MLE environment. A potential mapping is shown in the following snippet. See Example 4-6 for a complete definition of module mod object import mod.

```
CREATE OR REPLACE MLE ENV named_exports_env
   imports ('namedExports' MODULE named_exports_module);
/
```

```
CREATE OR REPLACE MLE MODULE mod_object_import_mod LANGUAGE JAVASCRIPT AS

import * as myMath from "namedExports";

function mySum() {...}
```

4.2 Export Functionality

Commonly exported identifiers in native JavaScript modules include variables, constants, functions, and classes.

Topics

Named Exports

The explicit use of identifiers in an export statement is referred to as using named exports in JavaScript.

Default Exports

As an alternative to named exports, a default export can be defined in JavaScript. A default export differs syntactically from a named export. Contrary to the latter, a default export does not require a set of curly brackets.

Private Identifiers

Any identifier not exported from a module is considered private and cannot be referenced outside the module's scope or in module call specifications.

4.2.1 Named Exports

The explicit use of identifiers in an export statement is referred to as using named exports in JavaScript.

Example 4-2 demonstrates the export of multiple functions using named exports.

Example 4-2 Function Export using Named Exports

This code snippet creates a module called <code>named_exports_module</code>, defines two functions, sum() and <code>difference()</code>, and then uses a named export to provide access for other modules to import the listed functions.

```
CREATE OR REPLACE MLE MODULE named_exports_module LANGUAGE JAVASCRIPT AS
function sum(a, b) {
   return a + b;
}
function difference(a, b) {
   return a - b;
}
export {sum, difference};
//
```



Make note of the $export\{\}$ statement at the end of the module. Named exports require the use of curly brackets when listing identifiers. Any identifier placed between the curly brackets is exported. Those not listed are not exported.

Rather than using the export statement towards the end of the module, it is also possible to prefix an identifier with the export keyword inline. The following example shows how the same module from the previous example can be rewritten with the export keyword provided inline with the JavaScript code.

Example 4-3 Function Export Using Export Keyword Inline

This code snippet creates a module called <code>inline_export_module</code> and defines two functions, <code>sum()</code> and <code>difference()</code>, which are both prefaced with the <code>export</code> keyword inline.

```
CREATE OR REPLACE MLE MODULE inline_export_module LANGUAGE JAVASCRIPT
AS

export function sum(a, b) {
   return a + b;
}

export function difference(a, b) {
   return a - b;
}
//
```

Both named_exports_module from Example 4-2 and inline_export_module are semantically identical. The method used to export the functions is the only syntactical difference.

4.2.2 Default Exports

As an alternative to named exports, a default export can be defined in JavaScript. A default export differs syntactically from a named export. Contrary to the latter, a default export does not require a set of curly brackets.



In line with the ECMAScript standard, only one default export is possible per module.

Example 4-4 Export a Class Using a Default Export

The following code snippet creates a module called default_export_module, defines a class called myMath, and defaults the class using a default export.

```
CREATE OR REPLACE MLE MODULE default_export_module
LANGUAGE JAVASCRIPT AS

export default class myMath {
   static sum(operand1, operand2) {
```



```
return operand1 + operand2;
}

static difference(operand1, operand2) {
    return operand1 - operand2;
}
}
```

4.2.3 Private Identifiers

Any identifier not exported from a module is considered private and cannot be referenced outside the module's scope or in module call specifications.

Example 4-5 Named Export of Single Function

The following code snippet creates a module called $private_export_module$, defines two functions, sum() and difference(), and exports the function sum() via named export. The function difference() is not included in the export statement, thus is only available within its own module's scope.

```
CREATE OR REPLACE MLE MODULE private_export_module
LANGUAGE JAVASCRIPT AS

function sum(a, b) {
   return a + b;
}

function difference(a, b) {
   return a - b;
}

export { sum };
//
```

4.3 Import Functionality

The import keyword allows developers to import functionality that has been exported by a source module.

Topics

Module Objects

A module object supplies a convenient way to import everything that has been exported by a module.

Named Imports

The ECMAScript standard specifies named imports. Rather than using an import name, you also have the option to specify identifiers.

Default Imports

Unlike named imports, default imports don't require the use of curly braces. This syntactical difference is also relevant to MLE's built-in modules.



4.3.1 Module Objects

A module object supplies a convenient way to import everything that has been exported by a module.

The module object provides a means to access all identifiers exported by a module and is used as a kind of namespace when referring to the imports.

Example 4-6 Module Object Definition

```
CREATE MLE ENV named_exports_env
    IMPORTS('namedExports' module named_exports_module);

CREATE OR REPLACE MLE MODULE mod_object_import_mod
LANGUAGE JAVASCRIPT AS

//the definition of a module object is demonstrated by the next line
import * as myMath from "namedExports"

function mySum() {
    const result = myMath.sum(4, 2);
    console.log(`the sum of 4 and 2 is ${result}`);
}

function myDifference() {
    const result = myMath.difference(4, 2);
    console.log(`the difference between 4 and 2 is ${result}`);
}

export {mySum, myDifference};
//
```

myMath identifies the module object and named_exports_module is the module name. Both sum() and difference() are available under the myMath scope in mod object import mod.

4.3.2 Named Imports

The ECMAScript standard specifies named imports. Rather than using an import name, you also have the option to specify identifiers.

Example 4-7 Named Imports Using Specified Identifiers

```
CREATE OR REPLACE MLE MODULE named_imports_module
LANGUAGE JAVASCRIPT AS

import {sum, difference} from "namedExports";

function mySum() {
   const result = sum(4, 2);
   console.log(`the sum of 4 and 2 is ${result}`);
}
```



```
function myDifference() {
    const result = difference(4, 2);
    console.log(`the difference between 4 and 2 is ${result}`);
}
export {mySum, myDifference};
/
```

Namespace clashes can ensue if multiple modules export the same identifier. To avoid this issue, you can provide an alias in the import statement.

Example 4-8 Named Imports with Aliases

```
CREATE OR REPLACE MLE MODULE named_imports_alias_module
LANGUAGE JAVASCRIPT AS

//note the use of aliases in the next line
import {sum as theSum, difference as theDifference} from "namedExports";

function mySum() {
   const result = theSum(4, 2);
   console.log(`the sum of 4 and 2 is ${result}`);
}

function myDifference() {
   const result = theDifference(4, 2);
   console.log(`the difference between 4 and 2 is ${result}`);
}

export {mySum, myDifference};
//
```

Instead of referencing the exported functions by their original names, the alias is used instead.

4.3.3 Default Imports

Unlike named imports, default imports don't require the use of curly braces. This syntactical difference is also relevant to MLE's built-in modules.

Example 4-9 Default Import

This example demonstrates the default import of myMathClass.

```
CREATE OR REPLACE MLE ENV default_export_env
    IMPORTS('defaultExportModule' MODULE default_export_module);

CREATE MLE MODULE default_import_module LANGUAGE JAVASCRIPT AS

//note the lack of curly braces in the next line
import myMathClass from "defaultExportModule";

export function mySum() {
    const result = myMathClass.sum(4, 2);
```



```
console.log(`the sum of 4 and 2 is ${result}`);
}
```

The same technique applies to the use of MLE's built-in modules such as the SQL driver. Example 4-10 demonstrates the use of the SQL driver in JavaScript code.

Example 4-10 Default Import with Built-in Module

```
CREATE MLE MODULE default_import_built_in_mod
LANGUAGE JAVASCRIPT AS

//note that there is no need to use MLE environments with built-in
modules
import oracledb from "mle-js-oracledb";

export function hello() {
    const options = {
        resultSet: false,
        outFormat: oracledb.OUT_FORMAT_OBJECT
    };
    const bindvars = [];

const conn = oracledb.defaultConnection();
const result = conn.execute('select user', bindvars, options);
console.log(`hello, ${result.rows[0].USER}`);
}
```

Unlike other examples using custom JavaScript modules, no MLE environment is required for importing a built-in module.



Server-Side JavaScript API Documentation for more information about the built-in JavaScript modules



MLE JavaScript Functions

This chapter introduces the use of call specifications to publish JavaScript functions so that they can be called from SQL and PL/SQL. MLE's support of OUT and IN OUT parameters is also discussed.

Topics

- Call Specifications for Functions
 Call specifications for modules and inlined MLE call specifications allow you to implement
 JavaScript functionality.
- OUT and IN OUT Parameters

5.1 Call Specifications for Functions

Call specifications for modules and inlined MLE call specifications allow you to implement JavaScript functionality.

Functions exported by an MLE JavaScript module can be published by creating call specifications. A JavaScript function published with a call specification can be called from anywhere a PL/SQL function or procedure can be called.

Alternatively, inlined MLE call specifications can be used to embed JavaScript code directly in the DDL. This option can be advantageous when you want to quickly implement a simple functionality using JavaScript.

Topics

- Creating a Call Specification for an MLE Module
 MLE call specification creation uses the generic CREATE FUNCTION RETURNS AS OF CREATE
 PROCEDURE AS Syntax, followed by MLE specific syntax.
- Creating an Inlined MLE Call Specification
 Inlined MLE call specifications embed JavaScript code directly in the CREATE FUNCTION and CREATE PROCEDURE DDLs.
- Choosing Inlined Versus Module MLE Call Specifications
 Each option provides its own advantages and disadvantages depending on your use case.
- Runtime Isolation for an MLE Call Specification
 - Dictionary Views for Call Specifications

 Metadata about JavaScript call specifications is available in the data dictionary using the
 [USER | ALL | DBA | CDB]_MLE_PROCEDURES views. The family of views maps call
 specifications (package, function, procedure) to JavaScript modules. This dictionary view
 is closely modeled after the * PROCEDURES views.



5.1.1 Creating a Call Specification for an MLE Module

MLE call specification creation uses the generic CREATE FUNCTION RETURNS AS OR CREATE PROCEDURE AS syntax, followed by MLE specific syntax.

Example 5-1 Creating MLE Call Specifications

This example walks you through the creation of an MLE module that exports two functions, and the creation of call specifications to publish those functions.

```
CREATE OR REPLACE MLE MODULE jsmodule
LANGUAGE JAVASCRIPT AS

export function greet(str) {
    console.log(`Hello, ${str}`);
  }
  export function concat(str1, str2) {
    return str1 + str2;
  }
/
```

The MLE module <code>jsmodule</code> exports two functions. The function <code>greet()</code> takes an input string argument and prints a simple greeting, while the function <code>concat()</code> takes two strings as input and returns the concatenated string as the result.

Because <code>greet()</code> does not return a value, you must create a PL/SQL procedure to publish it, as follows:

```
CREATE OR REPLACE PROCEDURE

GREET(str in VARCHAR2)

AS MLE MODULE jsmodule

SIGNATURE 'greet(string)';
```

The above call specification creates a PL/SQL procedure named <code>GREET()</code> in the schema of the current user. Running the procedure executes the exported function <code>greet()</code> in the JavaScript module <code>jsmodule</code>.

Note that it is not a requirement that the call specification has the same name (GREET) as the function being published (greet).

The MLE specific clause MLE MODULE <module name> specifies the JavaScript MLE module that exports the JavaScript function to be called.

The SIGNATURE clause specifies the name of the exported function to be called (in this case, greet), as well as, optionally, the list of argument types in parentheses. JavaScript MLE functions use TypeScript types in the SIGNATURE clause. In this example, the function accepts a JavaScript string. The PL/SQL VARCHAR2 string is converted to a JavaScript string before invoking the underlying JavaScript implementation. The SIGNATURE clause also allows the list of argument types to be omitted, in which case only the MLE function name is expected and MLE language types are inferred from the types given in the call specification's argument list.



The other exported function, concat(), can similarly be used to create a PL/SQL function:

```
CREATE OR REPLACE FUNCTION CONCATENATE(str1 in VARCHAR2, str2 in VARCHAR2)
   RETURN VARCHAR2
   AS MLE MODULE jsmodule
   SIGNATURE 'concat(string, string)';
/
```

The call specification in this case additionally specifies the PL/SQL return type of the created function. The value returned by the JavaScript function <code>concat()</code> (of type string) is converted to the type <code>VARCHAR2</code>.

The created procedure and function can be called as shown below with the result:

Topics

Components of an MLE Call Specification

The elements of an MLE call specification are listed along with descriptions.

MLE Module Clause

The MLE MODULE clause specifies the MLE module that exports the underlying JavaScript function for the call specification. The specified module must always be in the same schema as the call specification being created.

ENV Clause

The optional ENV clause specifies the MLE environment for module contexts in which this call specification will be executed.

SIGNATURE Clause

The SIGNATURE clause contains all the information necessary to map the MLE call specification to a particular function exported by the specified MLE module.



5.1.1.1 Components of an MLE Call Specification

The elements of an MLE call specification are listed along with descriptions.

Figure 5-1 MLE Call Specification Syntax

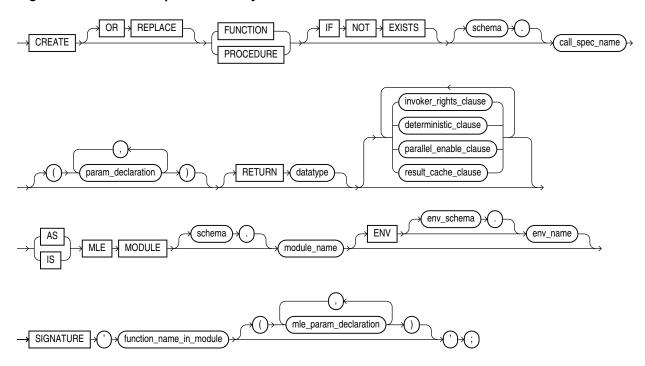


Table 5-1 Components of an MLE Call Specification

Element Name	Description
OR REPLACE	Specifies that the function should be replaced if it already exists. This clause can be used to change the definition of an existing function without dropping, recreating, and regranting object privileges previously granted on the function. Users who had previously been granted privileges on a recreated function or procedure can still access the function without being re-granted the privileges.
IF NOT EXISTS	Specifies that the function should be created if it does not already exist. If a function by the same name does exist, the statement is ignored without error and the original function body remains unchanged. Note that SQL*Plus will display the same output message regardless of whether the command is ignored or executed, ensuring that your DDL scripts remain idempotent.
	IF NOT EXISTS cannot be used in combination with ${\tt OR}$ REPLACE.
schema	Specifies the schema that will contain the call specification. If the schema is omitted, the call specification is created in the schema of the current user.



Table 5-1 (Cont.) Components of an MLE Call Specification

Element Name	Description
call_spec_name	Specifies the name of the call specification to be created. Call specifications are created in the default namespace, unlike MLE modules and environments, which use dedicated namespaces.
param_declaration	Specifies the call specification's parameters. If no parameters are specified, parentheses must be omitted.
RETURN datatype	Only used for functions and specifies the data type of the return value of the function. The return value can have any data type supported by MLE. Only the data type is specified; length, precision, or scale information must be omitted.
invoker_rights_clause	Specifies whether a function is invoker's or definer's rights. AUTHID CURRENT_USER creates an invoker's rights function or procedure. AUTHID DEFINER creates a definer's rights function or procedure. If the AUTHID clause is omitted, the call specification is created with definer's rights by default. The AUTHID clause on MLE call specifications has the exact same semantics as on PL/SQL functions and procedures.
deterministic_clause	Only used for functions and indicates that the function returns the same result value whenever it is called with the same values for its parameters. As with PL/SQL functions, this clause should not be used for functions that access the database in any way that might affect the return result of the function. The results of doing so will not be captured if the database chooses not to re-execute the function.

5.1.1.2 MLE Module Clause

The MLE MODULE clause specifies the MLE module that exports the underlying JavaScript function for the call specification. The specified module must always be in the same schema as the call specification being created.

An ORA-04103 error is thrown if the specified MLE module does not exist. Likewise, an ORA-01031 error is raised if the specified module is in a different schema from the created call specification.

5.1.1.3 ENV Clause

The optional ${\tt ENV}$ clause specifies the MLE environment for module contexts in which this call specification will be executed.

An ORA-04105 error is thrown if the specified environment schema object does not exist.

If this clause is omitted, the default environment is used. The default environment is simply an environment in its most basic state, with no module imports and no specified language options.



5.1.1.4 SIGNATURE Clause

The SIGNATURE clause contains all the information necessary to map the MLE call specification to a particular function exported by the specified MLE module.

Specifically, it includes two pieces of information:

- The name of the exported function in the specified MLE module.
- The MLE language parameter types (as opposed to the PL/SQL parameter types) for the function (Optional).

The SIGNATURE clause must be in the following form:

```
signature = path-spec "(" { language-type, "," } ")"
path-spec = import-spec | import-spec "." javascript-identifier
import-spec = javascript-identifier | "default"
```

The path specification describes the function to be called and can have the following two forms:

- A path specification can consist only of an import specification.
 - An import specification can be a JavaScript identifier that identifies a named export of the module, which must be a function. Alternatively, an import specification can be the reserved word, default. In this case, the default export of the module is used, which must be a function.
- A path specification can be a composite form consisting of an import specification, followed by a dot and a JavaScript identifier.
 - In this case, the import specification must refer to an object that has a property whose name matches the identifier listed after the dot. The value of the property needs to be a function.

The language-type can either be a built-in JavaScript type (e.g. string or number) or a type provided by MLE (e.g. OracleNumber or OracleDate) that is compatible with the corresponding PL/SQL argument. For more information about what types are provided by MLE through the built-in module mle-js-plsqltypes, see Server-Side JavaScript API Documentation.

function-name can include any alphanumeric characters as well as underscores and periods.

When the call specification is a function, the type of the return value is not specified in the SIGNATURE clause. Rather, the function can return any JavaScript type that is compatible with the PL/SQL type specified in the call specification's RETURN clause.



The parsing and resolution of the SIGNATURE clause happens lazily when the MLE function is executed for the first time. It is only at this point that any resolution or syntax errors in the SIGNATURE clause are reported, and not when the call specification is created.



Simplified SIGNATURE Clause

CREATE FUNCTION and CREATE PROCEDURE DDL statements also accept a simplified form of the SIGNATURE clause that only specifies the name of the exported function and leaves out the JavaScript language types of the parameters. The default PL/SQL-MLE language type mappings are used in this case.

This example demonstrates the creation of a call specification with a simplified SIGNATURE clause.

```
CREATE OR REPLACE FUNCTION concat
RETURN VARCHAR2
AS MLE MODULE jsmodule
SIGNATURE 'concat';
```

When the function concat is called from PL/SQL, the input VARCHAR2 parameters are converted to JavaScript string (the default type mapping for VARCHAR2) before calling the underlying JavaScript function.



MLE Type Conversions for more information about type mappings

5.1.2 Creating an Inlined MLE Call Specification

Inlined MLE call specifications embed JavaScript code directly in the CREATE FUNCTION and CREATE PROCEDURE DDLs.

If you want to quickly implement simple functionality using JavaScript, inlined MLE call specifications may be a good choice. With this option, you don't need to deploy a separate module containing the JavaScript code. Rather, the JavaScript function is built into the definition of the call specification itself.

The MLE LANGUAGE clause is used to specify that the function is implemented using JavaScript. The string following the language name is treated as the body of a JavaScript function that implements the functionality of the call specification. When the code is executed, PL/SQL parameters are automatically converted to the default JavaScript type and passed to the JavaScript function as parameters of the same name. Note that unquoted parameter names are mapped to all-uppercase JavaScript names. The value returned by a JavaScript function is converted to the return type of the PL/SQL call specification, just as with call specifications for MLE modules.



Oracle Database PL/SQL Packages and Types Reference for information about DBMS MLE subprograms for inlined call specifications

Example 5-2 Simple Inlined MLE Call Specification

The function created in the preceding inlined call specification results in the following when called from a SQL statement:

```
SELECT subWorld('Hello, World!');

SUBWORLD('HELLO, WORLD!')

Hello, Universe!
```

Example 5-3 Inlined MLE Call Specification Returning JSON

Note the use of double quotation marks in the function parameter name, strArgs, in Example 5-2. The inclusion preserves the parameter's letter case. Without quotation marks, the parameter name is mapped to an all-uppercase JavaScript name, as seen in this example.

```
CREATE OR REPLACE FUNCTION p string to JSON(inputString VARCHAR2)
RETURN JSON
AS MLE LANGUAGE JAVASCRIPT
q'~
  if ( INPUTSTRING === undefined ) {
    throw `must provide a string in the form of
key1=value1;...;keyN=valueN`;
  }
  let myObject = {};
  if ( INPUTSTRING.length === 0 ) {
   return myObject;
  const kvPairs = INPUTSTRING.split(";");
  kvPairs.forEach( pair => {
    const tuple = pair.split("=");
    if (tuple.length === 1) {
      tuple[1] = false;
    } else if ( tuple.length != 2 ) {
      throw "parse error: you need to use exactly one '=' between key
and value and not use '=' in either key or value";
   myObject[tuple[0]] = tuple[1];
  return myObject;
```



```
~';
/
```

The function created in the preceding inlined call specification results in the following when called from a SQL statement:

- Components of an Inlined MLE Call Specification
 The elements of an inlined MLE call specification are listed along with descriptions.
- Accessing Built-in Modules Using JavaScript Global Variables
 Rather than importing MLE built-in modules in the same way as call specifications for
 MLE modules, inlined MLE call specifications utilize prepopulated JavaScript globals to
 access built-in module functionality.

5.1.2.1 Components of an Inlined MLE Call Specification

The elements of an inlined MLE call specification are listed along with descriptions.

Figure 5-2 MLE Inlined Call Specification Syntax

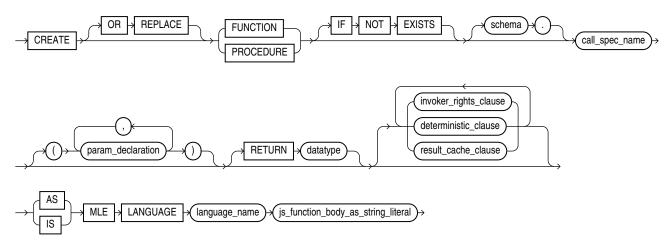


Table 5-2 Components of an Inlined MLE Call Specification

Element Name	Description
OR REPLACE	Specifies that the function should be replaced if it already exists. This clause can be used to change the definition of an existing function without dropping, recreating, and re-granting object privileges previously granted on the function. Users who had previously been granted privileges on a recreated function or procedure can still access the function without being re-granted the privileges.



Table 5-2 (Cont.) Components of an Inlined MLE Call Specification

Element Name	Description
IF NOT EXISTS	Specifies that the function should be created if it does not already exist. If a function by the same name does exist, the statement is ignored without error and the original function body remains unchanged. Note that SQL*Plus will display the same output message regardless of whether the command is ignored or executed, ensuring that your DDL scripts remain idempotent.
	IF NOT EXISTS cannot be used in combination with OR REPLACE.
schema	Specifies the schema that will contain the call specification. If the schema is omitted, the call specification is created in the schema of the current user.
call_spec_name	Specifies the name of the call specification to be created. Call specifications are created in the default namespace, unlike MLE modules and environments, which use dedicated namespaces.
param_declaration	Specifies the call specification's parameters. If no parameters are specified, parentheses must be omitted.
RETURN datatype	Only used for functions and specifies the data type of the return value of the function. The return value can have any data type supported by MLE. Only the data type is specified; length, precision, or scale information must be omitted.
invoker_rights_clause	Specifies whether a function is invoker's or definer's rights. AUTHID CURRENT_USER creates an invoker's rights function or procedure. AUTHID DEFINER creates a definer's rights function or procedure.
	If the AUTHID clause is omitted, the call specification is created with definer's rights by default. The AUTHID clause on MLE call specifications has the exact same semantics as on PL/SQL functions and procedures.
deterministic_clause	Only used for functions and indicates that the function returns the same result value whenever it is called with the same values for its parameters. As with PL/SQL functions, this clause should not be used for functions that access the database in any way that might affect the return result of the function. The results of doing so will not be captured if the database chooses not to reexecute the function.
MLE LANGUAGE	Specifies the language of the following code, for example, JavaScript. The string following the language name is interpreted as MLE language code implementing the desired functionality. For JavaScript, this embedded code is interpreted as the body of a JavaScript function.

5.1.2.2 Accessing Built-in Modules Using JavaScript Global Variables

Rather than importing MLE built-in modules in the same way as call specifications for MLE modules, inlined MLE call specifications utilize prepopulated JavaScript globals to access built-in module functionality.

Inlined MLE call specifications cannot import MLE modules, both built-in and custom. Instead, JavaScript global variables, such as the session variable, provide access to the functionality of built-in modules like the JavaScript MLE SQL driver. For more

information about the availability of objects in the global scope, see Working with the MLE JavaScript Driver.

If you need to use the functionality of a custom module, you must either use the import keyword with call specifications for MLE modules or use the async/await interface with dynamic MLE execution.



Server-Side JavaScript API Documentation for more information about the built-in JavaScript modules

5.1.3 Choosing Inlined Versus Module MLE Call Specifications

Each option provides its own advantages and disadvantages depending on your use case.

Inlined MLE call specifications can simplify the development workflow and provide a way to quickly implement simple JavaScript functionality, as there is no need to deploy a separate module containing the JavaScript code. This is a convenient option if you only need to implement a single JavaScript function. You can use JavaScript global variables to access the functionality of MLE built-in modules but, because inlined MLE call specifications are not associated with an MLE environment, modules cannot be imported directly.

Call specifications for MLE modules offer more flexibility in terms of complexity and ability to import functionality from other modules, built-in and custom. You also have the option to override the default JavaScript type mapping, which is not possible with MLE inlined call specifications. Call specifications for MLE modules should be used for larger pieces of JavaScript code as well as for code that you intend to reuse in other JavaScript code using imports.



Oracle Database PL/SQL Packages and Types Reference for information about DBMS MLE subprograms for MLE call specifications

5.1.4 Runtime Isolation for an MLE Call Specification

MLE uses execution contexts to maintain runtime state isolation. Call specifications are associated with separate contexts when they do not share the same user, module, and environment.

MLE execution contexts act as standalone, isolated runtime environments. All JavaScript code that shares an execution context has full access to all of its runtime state (e.g. any global variables previously defined). Otherwise, there is no way for code executing in one execution context to see or modify runtime state in another execution context. Execution contexts for call specifications are created transparently on the first call to any of the corresponding call specifications. For more information, see About MLE Execution Contexts.

When executing call specifications in a session, MLE loads the module specified in the call specification and calls the function(s) exported by that module. In order for the execution of

two call specifications to share the same execution context, they must export a function from the same MLE module, use the same environment, and be executed by the same user. SQL or PL/SQL calls on behalf of different users within the same session are never executed in the same execution context.

The runtime representation of a module is stateful. State constitutes, for example, variables in the JavaScript module as well as variables in the global scope accessible to code in the module. Within the same session, MLE may employ multiple module contexts to execute call specifications. If either the module or the environment referred to by a call specification change, any execution context is invalidated and an error is thrown. Example 5-4 demonstrates this concept.

Session state is very important for data integrity. Not catching errors related to changed session state (ORA-04106 for module changes and ORA-04107 for environment changes in JavaScript, as well as ORA-04068 for PL/SQL packages) can result in silent data corruption. Setting the initialization parameter

SESSION_EXIT_ON_PACKAGE_STATE_ERROR to TRUE forces sessions to be disconnected if the session state is invalidated. Because many applications capture session disconnect, this option can help simplify the recovery from the invalidation of existing session state. For more information about SESSION_EXIT_ON_PACKAGE_STATE_ERROR, see *Oracle Database Reference*.



Storing state in packages and JavaScript modules is not recommended. Session state is best handled by the database.

All definer's rights call specifications that publish functions from the same MLE module (and use the same environment) will share the same execution context because all execution happens on behalf of the definer. Conversely, there is a separate execution context per calling user when a call specification is declared as invoker's rights.

For more information about how to build a call specification, see Components of an MLE Call Specification.

See Also:

Oracle Database PL/SQL Language Reference for information about using SESSION_EXIT_ON_PACKAGE_STATE_ERROR to specify behavior when PL/SQL package state is invalidated

Example 5-4 Execution Context Dependencies

This example demonstrates the fact that if a module or environment changes, any associated execution context(s) are invalidated.

CREATE OR REPLACE MLE MODULE count_module LANGUAGE JAVASCRIPT AS

let myCounter = 0;



Another user invoking the function from a different session, we'll say session 2, creates another execution context, separate from the first session's context:

The user in session 1 recreates the MLE module with some new comments added to the function:

```
CREATE OR REPLACE MLE MODULE count_module
LANGUAGE JAVASCRIPT AS

let myCounter = 0;

/**
 * increments a counter before returning the value
 * to the caller
 *@returns {number} the value of the counter
 */
export function incrementCounter() {
    return ++myCounter;
}
//
```



This operation signals to all execution contexts referring to <code>count_module</code> that their session state should be invalidated. Session 2 gets an error in response to the invalidation:

```
SQL> SELECT increment_and_get_counter;

SELECT increment_and_get_counter

*

ERROR at line 1:

ORA-04106: Module USER2.COUNT_MODULE referred to by

INCREMENT_AND_GET_COUNTER has been modified since the execution context was created.
```

The next invocation of the function in session 2 starts off with a reinitialized session state:

```
SQL> SELECT increment_and_get_counter;

INCREMENT_AND_GET_COUNTER

1
```

Just as with PL/SQL packages, invoking the function from session 1 does not result in an error. Nevertheless, the session state has been discarded as shown by a subsequent call to the function:

If the initialization parameter <code>SESSION_EXIT_ON_PACKAGE_STATE_ERROR</code> is set to <code>TRUE</code> in session 2, the <code>ORA-04106</code> error is thrown and the connection to the database is cut:

```
ALTER SESSION SET SESSION_EXIT_ON_PACKAGE_STATE_ERROR = TRUE; SELECT increment and get counter;
```

Result:

```
SELECT increment_and_get_counter

*

ERROR at line 1:

ORA-04106: Module USER2.COUNT_MODULE referred to by

INCREMENT_AND_GET_COUNTER has been modified since the execution context was created.

ERROR:

ORA-03114: not connected to ORACLE
```



5.1.5 Dictionary Views for Call Specifications

Metadata about JavaScript call specifications is available in the data dictionary using the $[USER \mid ALL \mid DBA \mid CDB]$ _MLE_PROCEDURES views. The family of views maps call specifications (package, function, procedure) to JavaScript modules. This dictionary view is closely modeled after the * PROCEDURES views.

For more information about * MLE PROCEDURES, see Oracle Database Reference.

Example 5-5 Show JavaScript Call Specification Metadata

```
SELECT OBJECT_NAME, PROCEDURE_NAME, SIGNATURE, ENV_NAME, MODULE_NAME FROM USER MLE PROCEDURES;
```

SQL*Plus output:

OBJECT_NAME MODULE_NAME	PROCEDURE_NAME	SIGNATURE	ENV_NAME	
CONCATENATE		<pre>concat(string, string)</pre>		JSMODULE
DO_NOTHING		doNothing(string)		JSMODULE

5.2 OUT and IN OUT Parameters

Use out and in out parameters with MLE JavaScript functions.

MLE JavaScript functions support IN OUT and OUT parameters in addition to IN parameters, just as they are supported in PL/SQL functions and procedures. These are declared as IN OUT and OUT in the list of arguments of an MLE call specification.

Because JavaScript has no notion of output parameters, the JavaScript implementation instead accepts objects that wrap the parameter value. Concretely, the shape of these wrapper objects is described by the following generic TypeScript interfaces InOut and Out (for IN OUT and OUT parameters, respectively):

```
Interface InOut<T> {
     Value : T;
}

Interface Out<T> {
     Value : T;
}
```

Note that OUT and IN OUT parameters are passed to JavaScript functions as JavaScript objects whose only property, value, exposes the value of the argument. This means that, in order to read, write, and use the value of an OUT or IN OUT argument, it must first be unwrapped by accessing its value property. This is done in order to simulate a pass-by-reference implementation, which does not exist in JavaScript. For example, the substitute () function in Example 5-6 must first unwrap its IN OUT argument, sentence, by retrieving its

value property before calling match() on it. Attempting to call match() on sentence directly would fail, as sentence is only the value wrapper. These wrapper classes are never needed in DBMS MLE, which does not make use of OUT and IN OUT parameters.

Example 5-6 OUT and IN OUT Parameters with JavaScript

Consider an MLE function, substitute(), that takes a VARCHAR2 IN OUT parameter, sentence, and replaces all occurrences of the second parameter, replaceThis, with the third parameter, withThat, then returns the number of occurrences of replaceThis in sentence.

```
CREATE OR REPLACE MLE MODULE in out example mod
LANGUAGE JAVASCRIPT AS
export function substitute (sentence, replaceThis, withThat) {
     * substitute: substitutes `replaceThis` in `sentence` with
                    `replaceThat`
     * parameters:
     * - sentence: the input sentence
     * - replaceThis: a word to be replaced in `sentence`
     * - withThat: the new word to be used instead of `replaceThis`
     */
    const occurrences =
        (sentence.value.match(replaceThis) || []).length;
    sentence.value = sentence.value.replace(replaceThis, withThat);
    return occurrences;
CREATE OR REPLACE FUNCTION f substitute(
   p sentence IN OUT VARCHAR2,
   p_replaceThis IN VARCHAR2, p withThat IN VARCHAR2
RETURN NUMBER
AS MLE MODULE in out example mod
SIGNATURE 'substitute(InOut<string>, string, string)';
```

The SIGNATURE clause of the call specification lists the parameter type of the JavaScript function's sentence parameter as InOut<string>. The input VARCHAR2 value is therefore converted to a JavaScript string, that is then wrapped in an object and passed to the JavaScript function substitute().



```
l_occurrences pls_integer;
BEGIN

dbms_output.put_line('sentence before: ' || l_sentence);
l_occurrences := f_substitute(
    l_sentence, l_replaceThis, l_withThat);
if l_occurrences <> 0 then
    dbms_output.put_line('sentence after: ' || l_sentence);
else
    dbms_output.put_line('no text replacement performed');
end if;
END;
/
Result:
sentence before: people are enjoying the rain
sentence after: people are enjoying the sun
```



6

Calling PL/SQL and SQL from the MLE JavaScript SQL Driver

- Introduction to the MLE JavaScript SQL Driver
- Selecting Data Using the MLE JavaScript Driver
- Data Modification
- Bind Variables
- PL/SQL Invocation from the MLE JavaScript SQL Driver
- Error Handling in SQL Statements
- Working with JSON Data
 The use of JSON data as part of a relational structure, more specifically the use of JSON columns in (relational) tables, is described.
- Using Large Objects (LOB) with MLE
- API Differences Between node-oracledb and mle-js-oracledb
 There are several differences between node-oracledb and mle-js-oracledb, including the methods for handling connection management and type mapping.

6.1 Introduction to the MLE JavaScript SQL Driver

The MLE JavaScript driver is closely modeled after the client-side Oracle SQL driver for Node.js, node-oracledb.

This close relationship between the server-side and client-side drivers reduces the effort required for porting client-side JavaScript code from Node.js or Deno to the database. Functionality that cannot be reasonably mapped to the server-side environment is omitted from MLE and the MLE JavaScript driver and will throw errors.

This helps you identify those parts of the code requiring changes. Furthermore, the MLE JavaScript implementation is a pure JavaScript implementation. Certain features not part of the ECMAScript standard are not available in MLE such as the window object or direct file and network I/O.

The mle-js-oracledb SQL driver defaults to a synchronous operating model and partially supports asynchronous execution via async/await.



Production code should adhere to industry best practices for error handling and logging, both of which have been omitted from examples in this chapter for the sake of clarity. Additionally, most of the examples feature the synchronous execution model due to its greater readability.

See Also:

- API Differences Between node-oracledb and mle-js-oracledb
- Server-Side JavaScript API Documentation for more information about the built-in JavaScript modules

Topics

- Working with the MLE JavaScript Driver
 Generic workflow for working with the MLE JavaScript driver.
- Connection Management in the MLE JavaScript Driver
- Introduction to Executing SQL Statements
- Processing Comparison Between node-oracledb and mle-js-oracledb
 The node-oracledb documentation recommends the use of the async/await interface. Due to the nature of client-server interactions, most of the processing involved between node and the database is executed asynchronously.

6.1.1 Working with the MLE JavaScript Driver

Generic workflow for working with the MLE JavaScript driver.

At a high level, working with the MLE JavaScript driver is very similar to using the client-side node-oracledb driver, namely:

- 1. Get a connection handle to the existing database session.
- 2. Use the connection to execute a SQL statement.
- 3. Check the result object returned by the statement executed, as well as any database errors that may have occurred.
- 4. In the case of select statements, iterate over the resulting cursor.
- For statements manipulating data, decide whether to commit or roll the transaction back.

Applications that aren't ported from client-side Node.js or Deno can benefit from coding aids available in the MLE JavaScript SQL driver, such as a number of frequently used variables that are available in the global scope. These variables include the following:

- oracledb for the OracleDb driver object
- session for the default connection object
- soda for the SodaDatabase object

Additionally, the following types are available:

- OracleNumber
- OracleClob
- OracleBlob
- OracleTimestamp



- OracleTimestampTZ
- OracleDate

The availability of these objects in the global scope reduces the need for writing boilerplate code. For details about global symbols available with MLE, see npmjs.com.

6.1.2 Connection Management in the MLE JavaScript Driver

Considerations when dealing with connection management in the MLE JavaScript driver. Connection management in the MLE JavaScript driver is greatly simplified. Because a database session will already exist by the time a stored procedure is invoked, you don't need to worry about establishing and tearing down connections, connection pools, and secure credential management, to name just a few.

You need only be concerned with the <code>getDefaultConnection()</code> method from the <code>mle-js-oracledb</code> module or use the global session object.

6.1.3 Introduction to Executing SQL Statements

A single SQL or PL/SQL statement can be executed by the Connection class's execute() method. Query results can either be returned in a single JavaScript array or fetched in batches via a ResultSet object.

Fetching as ResultSet offers more control over the fetch operation whereas using arrays requires fewer lines of code and provides performance benefits.

The JavaScript MLE SQL driver allows you to shorten the code demonstrated by Example 6-1 by accessing global variables, shown in Example 6-2. For more information about global variables available with MLE, see Working with the MLE JavaScript Driver.

Example 6-1 Getting Started with the MLE JavaScript SQL Driver

The following code demonstrates how to import the MLE JavaScript SQL driver into the current module's namespace.



```
{
           outFormat: oracledb.OUT FORMAT OBJECT
    );
    if (result.rows.length > 0) {
        for (let row of result.rows) {
            console.log(`The query found a row:
               manager id: ${row.MANAGER ID}
                department id: ${row.DEPARTMENT ID}
               department name: ${row.DEPARTMENT NAME}`);
    } else {
        console.log(`no data found for manager ID ${managerID}`);
  } catch (err) {
   console.error(`an error occurred while processing the query: $
{err.message}`);
}
export { queryExample };
```

The only function in the module, <code>queryExample()</code>, <code>selects</code> a single row from the HR departments table using a bind variable by calling <code>connection.execute()</code>. The value of the bind variable is passed as a parameter to the function. Another parameter passed to <code>connection.execute()</code> indicates that each row returned by the query should be provided as a JavaScript object.

If data has been found for a given managerID, it is printed on the screen. By default, the call to <code>console.log()</code> is redirected to <code>DBMS_OUTPUT</code>. Should there be no rows returned a message indicating this fact is printed on the console.

The call specification in the following snippet allows the code to be invoked in the database.

```
CREATE OR REPLACE PROCEDURE p_js_sql_query_ex(
    p_manager_id number)
AS MLE MODULE js_sql_mod
SIGNATURE 'queryExample(number)';
//
```

Provided the defaults are still in place, invoking $p_js_sql_query_ex$ displays the following:

```
SQL> EXEC p_js_sql_query_ex(103)
The query found a row:
manager_id: 103
department_id: 60
department name: IT
```



See Also:

Server-Side JavaScript API Documentation for more information about the built-in JavaScript modules, including mle-js-oracledb

Example 6-2 Use Global Variables to Simplify SQL Execution

```
CREATE OR REPLACE MLE MODULE intro mod globals LANGUAGE JAVASCRIPT AS
function queryExample(managerID){
    if (managerID === undefined) {
        throw "Parameter managerID has not been provided to queryExample()";
    const result = session.execute(`
        SELECT manager id, department id, department name
        FROM hr.departments
        WHERE manager id = :id`,
            managerID
        ],
        {
            outFormat: oracledb.OUT FORMAT OBJECT
    if(result.rows.length > 0){
        for(let row of result.rows) {
            console.log(
                `The query found a row:
                manager id: ${row.MANAGER ID}
                department id: ${row.DEPARTMENT ID}
                department name: ${row.DEPARTMENT NAME}`
            );
    } else {
        console.log(`no data found for manager ID ${managerID}`);
}
export{queryExample};
```

The call specification in the following snippet allows the code to be invoked in the database.

```
CREATE OR REPLACE PROCEDURE intro_mod_globals_query_ex(
    p_manager_id NUMBER)
AS MLE MODULE intro_mod_globals
SIGNATURE 'queryExample(number)';
//
```



```
intro_mod_globals_query_ex

SQL> BEGIN
    intro_mod_globals_query_ex(100);
END;
/

The query found a row:
manager_id: 100
department_id: 90
department name: Executive
```

6.1.4 Processing Comparison Between node-oracledb and mle-js-oracledb

The node-oracledb documentation recommends the use of the async/await interface. Due to the nature of client-server interactions, most of the processing involved between node and the database is executed asynchronously.

The MLE JavaScript driver does not require asynchronous processing. Like the PL/SQL driver, this is thanks to the driver's location within the database. The MLE JavaScript driver understands the async/await syntax, however, it processes requests synchronously under the hood.

Unlike the node-oracledb driver, the MLE JavaScript SQL driver returns rows as objects (oracledb.OUT_FORMAT_OBJECT) rather than arrays (oracledb.OUTFORMAT_ARRAY) when using the ECMAScript 2022 syntax. Code still relying on the deprecated require syntax remains backwards compatible by returning rows as an array.



A promise-based interface is not provided with the MLE JavaScript driver.

6.2 Selecting Data Using the MLE JavaScript Driver

Data can be selected using Direct Fetches or ResultSet objects. You can choose between arrays and objects as the output format. The default is to return data as Direct Fetch using JavaScript objects.

Topics

- Direct Fetch: Arrays
- Direct Fetch: Objects
- Fetching Rows as ResultSets: Arrays
- Fetching Rows as ResultSets: Iterating Over ResultSet Objects



6.2.1 Direct Fetch: Arrays

Direct Fetches are the default in the MLE JavaScript driver.

Direct Fetches provide query results in result.rows. This is a multidimensional JavaScript array if you specify the outFormat as oracledb.OUT_FORMAT_ARRAY. Iterating over the rows allows you to access columns based on their position in the select statement. Changing the column order in the select statement requires modifications in the parsing of the output. Because this can lead to bugs that are hard to detect, the MLE JavaScript driver returns objects by default (oracledb.OUT_FORMAT_OBJECT), rather than arrays.

Example 6-3 demonstrates the use of Direct Fetches using the synchronous execution model.

Example 6-3 Selecting Data Using Direct Fetch: Arrays

```
CREATE OR REPLACE MLE MODULE dir fetch arr mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function directFetch() {
    const conn = oracledb.defaultConnection();
    const result = conn.execute(
        `SELECT department id, department name
        FROM hr.departments`,
        {outFormat: oracledb.OUT FORMAT ARRAY}
    );
    for (let row of result.rows) {
        const deptID = String(row[0]).padStart(3, '0');
        const deptName = row[1];
        console.log(`department ID: ${deptID} - department name: $
{deptName}`);
    }
CREATE OR REPLACE PROCEDURE dir fetch arr proc
   AS MLE MODULE dir fetch arr mod
    SIGNATURE 'directFetch()';
EXEC dir fetch arr_proc();
Result:
department ID: 010 - department name: Administration
department ID: 020 - department name: Marketing
department ID: 030 - department name: Purchasing
department ID: 270 - department name: Payroll
```

The execute() function returns a result object. Depending on the statement type (select, insert, delete, etc) different properties are available for further processing.

For information about mle-js-oracledb, see Server-Side JavaScript API Documentation.

6.2.2 Direct Fetch: Objects

JavaScript objects are returned by default when using Direct Fetch. To address potential problems with the ordering of columns in the select list, results are returned as JavaScript objects rather than as arrays.

Example 6-4 Selecting Data Using Direct Fetch: Objects

```
CREATE OR REPLACE MLE MODULE dir fetch obj mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function directFetchAsObject() {
    const conn = oracledb.defaultConnection();
    const result = conn.execute(
        `SELECT department id, department name
        FROM hr.departments`,
        { outFormat: oracledb.OBJECT }
        // Note that oracledb.OBJECT is the default, outFormat is
defined
        // here for clarity
   );
    for (let row of result.rows) {
        const deptID = String(row.DEPARTMENT ID).padStart(3, '0');
        const deptName = row.DEPARTMENT NAME;
        console.log(`department ID: ${deptID} - department name: $
{deptName}`);
    }
}
CREATE OR REPLACE PROCEDURE dir fetch obj proc
   AS MLE MODULE dir fetch obj mod
    SIGNATURE 'directFetchAsObject()';
EXEC dir fetch obj proc();
Result:
department ID: 010 - department name: Administration
department ID: 020 - department name: Marketing
department ID: 030 - department name: Purchasing
department ID: 270 - department name: Payroll
```

Unlike PL/SQL, JavaScript doesn't support the concept of named parameters. The <code>execute()</code> method accepts the SQL statement, <code>bindParams</code>, and options, in that exact order. The query doesn't use bind variables, thus an empty array matches the function's signature.

The third parameter in this example instructs the MLE JavaScript driver to return rows as an object. Instead of accessing columns positionally, it is possible to use column names instead. This approach offers greater resilience to changes in the query text.



Server-Side JavaScript API Documentation for more information about the mle-js-oracledb built-in module

6.2.3 Fetching Rows as ResultSets: Arrays

You can use <code>ResultSet</code> objects as an alternative to using <code>Direct Fetches</code>. In addition to using <code>Direct Fetches</code>, it is possible to use <code>ResultSet</code> objects. A <code>ResultSet</code> is created when the option property <code>resultSet</code> is set to true. <code>ResultSet</code> rows can be fetched using <code>getRow()</code> or <code>getRows()</code>.

Because rows are fetched as JavaScript objects by default instead of as arrays, outFormat must be defined as oracledb.OUT FORMAT ARRAY in order to fetch rows as a ResultSet.

Example 6-5 Fetching Rows Using a ResultSet

```
CREATE OR REPLACE MLE MODULE resultSet mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function resultSetFetch() {
    const conn = oracledb.defaultConnection();
    const result = conn.execute(
        `SELECT department id, department name
         FROM hr.departments`,
        { resultSet: true, outFormat: oracledb.OUT FORMAT ARRAY }
    );
    const rs = result.resultSet;
    let row;
    while ((row = rs.getRow())) {
        const deptID = String(row[0]).padStart(3, '0');
        const deptName = row[1];
        console.log(`department ID: ${deptID} - department name: $
{deptName}`);
    }
    rs.close();
```

Note that the fetch operation specifically requested an array rather than an object. Objects are returned by default.

```
CREATE OR REPLACE PROCEDURE dir_fetch_arr_proc
    AS MLE MODULE dir_fetch_arr_mod
    SIGNATURE 'directFetch()';
/
```



```
EXEC dir_fetch_arr_proc();

Result:

department ID: 010 - department name: Administration department ID: 020 - department name: Marketing department ID: 030 - department name: Purchasing
```

department ID: 270 - department name: Payroll

6.2.4 Fetching Rows as ResultSets: Iterating Over ResultSet Objects

In addition to the <code>ResultSet.getRow()</code> and <code>ResultSet.getRows()</code> functions, the MLE JavaScript driver's <code>ResultSet</code> implements the iterable and iterator protocols, simplifying the process for iterating over the <code>ResultSet</code>.

Using either the iterable or iterator protocols is possible. Both greatly simplify working with ResultSets. Each option is demonstrated in Example 6-6 and Example 6-7.



ResultSet objects must be closed once they are no longer needed.

Example 6-6 Using the Iterable Protocol with ResultSets

This example shows how to use the iterable protocol as an alternative to ResultSet.getRow(). Rather than providing an array of column values, the JavaScript objects are returned instead.

```
CREATE OR REPLACE MLE MODULE rs iterable mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function resultSetFetchAsObject() {
    const conn = oracledb.defaultConnection();
    const result = conn.execute(
        `SELECT department id, department name
        FROM hr.departments`,
        [],
            resultSet: true,
            outFormat: oracledb.OUT FORMAT OBJECT
    );
    const rs = result.resultSet;
    for (let row of rs) {
        const deptID = String(row.DEPARTMENT ID).padStart(3, '0');
        const deptName = row.DEPARTMENT NAME;
        console.log(`department ID: ${deptID} - department name: $
{deptName}`);
    }
    rs.close();
}
```



```
CREATE OR REPLACE PROCEDURE rs_iterable_proc

AS MLE MODULE rs_iterable_mod

SIGNATURE 'resultSetFetchAsObject()';

EXEC rs_iterable_proc();

Result:

department ID: 010 - department name: Administration department ID: 020 - department name: Marketing department ID: 030 - department name: Purchasing
...

department ID: 270 - department name: Payroll
```

Example 6-7 Using the Iterator Protocol with ResultSets

This example shows how you can iterate over a ResultSet using the iterator protocol. This is a reimplementation of Example 6-6.

```
CREATE OR REPLACE MLE MODULE rs iterator mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function resultSetFetchAsObjectIterator() {
    const conn = oracledb.defaultConnection();
    const result = conn.execute(
        `SELECT department id, department name
         FROM hr.departments`,
        [],
            resultSet: true,
            outFormat: oracledb.OUT FORMAT OBJECT
    );
    const iterator = result.resultSet.iterator();
    let row = iterator.next();
    while (! row.done) {
        const deptID = String(row.value.DEPARTMENT ID).padStart(3, '0');
        const deptName = row.value.DEPARTMENT NAME;
        console.log(`department ID: ${deptID} - department name: $
{deptName}`);
        row = iterator.next();
    result.resultSet.close();
}
CREATE OR REPLACE PROCEDURE rs iterator proc
    AS MLE MODULE rs iterator mod
    SIGNATURE 'resultSetFetchAsObjectIterator()';
```



```
EXEC rs_iterator_proc();

Result:

department ID: 010 - department name: Administration
department ID: 020 - department name: Marketing
department ID: 030 - department name: Purchasing
...
department ID: 270 - department name: Payroll
```

6.3 Data Modification

Modify data using the MLE JavaScript driver.

In addition to selecting data, it is possible to insert, update, delete, and merge data using the MLE JavaScript driver. The same general workflow can be applied to these operations as you would use when selecting data.

Example 6-8 Updating a Row Using the MLE JavaScript SQL Driver

The result object's rowsAffected property can be interrogated to determine how many rows have been affected by the update. The JavaScript function updateCommissionExampleEmpID145() returns the number of rows affected to the caller. In this instance, the function will return 1.

An alternative method to update data is to use the connection.executeMany() method. This function works best when used with bind variables.

6.4 Bind Variables

Use bind variables to control data that is passed into or retrieved from the database. SQL and PL/SQL statements may contain bind variables, indicated by colon-prefixed identifiers. These parameters indicate where separately specified values are substituted in a statement when it is executed, or where values are to be returned after execution.

Three different kinds of bind variables exist in the Oracle database:

IN bind variables



- OUT bind variables
- IN OUT bind variables

IN binds are values passed into the database. OUT binds are used to retrieve data from the database. IN OUT binds are passed in and may return a different value after the statement executes.

Using bind variables is recommended in favor of constructing SQL or PL/SQL statements through string concatenation or template literals. Both performance and security can benefit from the use of bind variables. When bind variables are used, the Oracle database does not have to perform a resource and time consuming hard-parse operation. Rather, it can reuse the cursor already present in the cursor cache instead.



Bind variables cannot be used in DDL statements such as CREATE TABLE, and neither can they substitute text of a query, only data.

Topics

- Using Bind-by-Name vs Bind-by-Position
 Bind variables are used in two ways: bind by name and bind by position. You must pick one for a given SQL command as the options are mutually exclusive.
- RETURNING INTO Clause
- Batch Operations

6.4.1 Using Bind-by-Name vs Bind-by-Position

Bind variables are used in two ways: bind by name and bind by position. You must pick one for a given SQL command as the options are mutually exclusive.

Topics

- Named Bind Variables
- Positional Bind Variables

6.4.1.1 Named Bind Variables

Binding by name requires the bind variable to be a string literal, prefixed by a colon. In the case of named binds, the bindParams argument to the connection.execute() function should ideally be provided with the following properties of each bind variable defined.

Property	Description
dir	The bind variable direction
val	The value to be passed to the SQL statement
type	The data type



Example 6-9 Using Named Bind Variables

```
CREATE OR REPLACE MLE MODULE named bind mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function bindByName(deptName, sal) {
    const conn = oracledb.defaultConnection();
    const result = conn.execute(
        `SELECT e.first name || ' ' || e.last name employee name,
e.salarv
         FROM employees e LEFT JOIN departments d
         ON (e.department id = d.department id)
         WHERE nvl(d.department name, 'n/a') = :deptName
         AND salary > :sal
         ORDER BY e.employee_id`,
            deptName: {
                dir: oracledb.BIND IN,
                val: deptName,
                type: oracledb.STRING
            },
            sal: {
                dir: oracledb.BIND IN,
                val: sal,
                type: oracledb.NUMBER
            }
        },
        { outFormat : oracledb.OUT FORMAT OBJECT }
    );
    console.log(`Listing employees working in ${deptName} with a
salary > ${sal}`);
    for (let row of result.rows) {
        console.log(`${row.EMPLOYEE NAME.padEnd(25)} - $
{row.SALARY}`);
    }
}
```

The bindParams argument to connection.execute() defines two named bind parameters:

- deptName
- sal

In this example, the function's input parameters match the names of the bind variables, which improves readability but isn't a requirement. You are free to assign bind variable names as long as the mapping in bindParams is correct.

6.4.1.2 Positional Bind Variables

Instead of using named bind parameters, you can alternatively provide bind-variable information as an array.

The number of elements in the array must match the number of bind parameters in the SQL text. Rather than mapping by name, the mapping of bind variable and value is



based on the position of the bind variable in the text and position of the item in the bind array.

Example 6-10 Using Positional Bind Variables

This example demonstrates the use of positional bind variables and represents a reimplementation of Example 6-9

```
CREATE OR REPLACE MLE MODULE positional bind mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function bindByPosition(deptName, sal) {
    const conn = oracledb.defaultConnection();
    const result = conn.execute(
        `SELECT e.first name || ' ' || e.last name employee name, e.salary
         FROM employees e LEFT JOIN departments d
         ON (e.department id = d.department id)
         WHERE nvl(d.department name, 'n/a') = :deptName
         AND salary > :sal
         ORDER BY e.employee id`,
        [
            deptName,
            sal
        ],
        { outFormat : oracledb.OUT FORMAT OBJECT }
    console.log(`Listing employees working in ${deptName} with a salary > $
{sal}`);
    for (let row of result.rows) {
        console.log(`${row.EMPLOYEE NAME.padEnd(25)} - ${row.SALARY}`);
    }
}
```

In this example, bindParams is an array rather than an object. The mapping between bind variables in the SQL text to values is done by position. The first item in the bindParams array maps to the first occurrence of a placeholder in the SQL text and so on.

6.4.2 RETURNING INTO Clause

The use of the RETURNING INTO clause is described.

The RETURNING INTO clause allows you to

- Fetch values changed during an update
- Return auto-generated keys during a single-row insert operation
- List rows deleted

Example 6-11 Using the RETURNING INTO Clause

This example shows how to retrieve the old and new values after an update operation. These values can be used for further processing.

```
CREATE OR REPLACE MLE MODULE ret_into_mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function updateReturningIntoOldNew(firstEmpID, lastEmpID) {
   const conn = oracledb.defaultConnection();
```



```
const result = conn.execute(
        `UPDATE employees
         SET last name = upper(last name)
         WHERE employee id between :firstEmpID and :lastEmpID
         RETURNING old last name, new last name
         INTO :oldLastName, :newLastName`,
            firstEmpID: {
                dir: oracledb.BIND IN,
                val: firstEmpID,
                type: oracledb.NUMBER
            } ,
            lastEmpID: {
                dir: oracledb.BIND IN,
                val: lastEmpID,
                type: oracledb.NUMBER
            },
            oldLastName: {
                type: oracledb.STRING,
                dir: oracledb.BIND OUT
            },
            newLastName: {
                type: oracledb.STRING,
                dir: oracledb.BIND OUT
            }
   );
   console.log(
        `update() completed successfully:
        - old values: ${JSON.stringify(result.outBinds.oldLastName)}
        - new values: ${JSON.stringify(result.outBinds.newLastName)}`
   );
}
```

This example features both IN and OUT bind variables:

- firstEmpID and lastEmpID specify the data range to be updated
- oldLastName is an array containing all the last names as they were before the update
- newLastName is another array containing the new values

6.4.3 Batch Operations

In addition to calling the connection.execute() function, it is possible to use connection.executeMany() to perform batch operations.

Using <code>connection.executeMany()</code> is like calling <code>connection.execute()</code> multiple times but requires less work. This is an efficient way to handle batch changes, for example, when inserting or updating multiple rows. The <code>connection.executeMany()</code> method cannot be used for queries.



connection.execute() expects an array containing variables to process by the SQL statement. The bindData array in Example 6-12 contains multiple JavaScript objects, one for each bind variable defined in the SQL statement. The for loop constructs the objects and adds them to the bindData array.

In addition to the values to be passed to the batch operation, the MLE JavaScript SQL driver needs to know about the values' data types. This information is passed as the bindDefs property in the connection.executeMany() options parameter. Both old and new last names in Example 6-12 are character strings with the changeDate defined as a date.

Just as with the <code>connection.execute()</code> function, <code>connection.executeMany()</code> returns the <code>rowsAffected</code> property, allowing you to quickly identify how many rows have been batch processed.

Example 6-12 Performing a Batch Operation

This example extends Example 6-10 by inserting the old and new last names into an audit table.

```
CREATE OR REPLACE MLE MODULE batch mod
LANGUAGE JAVASCRIPT AS
export function updateReturningIntoOldNew(firstEmpID, lastEmpID) {
    // convert the last name of a range of employees to UPPER case ...
    let result = conn.execute(...);
    let bindData = [];
    const changeDate = new Date();
    for (let i = 0; i < result.outBinds.oldLastName.length; i++) {</pre>
        bindData.push(
            {
                oldLastName: result.outBinds.oldLastName[i],
                newLastName: result.outBinds.newLastName[i],
                changeDate : changeDate
        )
    }
    // ... and leave an audit trail
    result = conn.executeMany(
        `insert into EMPLOYEES AUDIT OPERATIONS (
            old last name,
            new last name,
            change date
        ) values (
            :oldLastName,
            :newLastName,
            :changeDate
  )`,
        bindData,
        {
            bindDefs: {
                oldLastName: { type: oracledb.STRING, maxSize: 30 },
                newLastName: { type: oracledb.STRING, maxSize: 30 },
                changeDate : { type: oracledb.DATE }
            }
    );
```



```
return result.rowsAffected;
}
/
```

After the initial update statement completes, the database provides the old and new values of the <code>last_name</code> column affected by the update in the <code>result</code> object's <code>outBinds</code> property. Both <code>oldLastName</code> and <code>newLastName</code> are arrays. The array length represents the number of rows updated.

6.5 PL/SQL Invocation from the MLE JavaScript SQL Driver

Use the MLE JavaScript driver to call functions and procedures from PL/SQL. The vast majority of the Oracle Database's API is provided in PL/SQL. As such, PL/SQL and JavaScript are able to co-exist in the database. Calling PL/SQL from the MLE JavaScript driver is very similar to calling SQL statements.

Example 6-13 Calling PL/SQL from JavaScript

```
CREATE OR REPLACE MLE MODULE plsql js mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function plsqlExample() {
    const conn = oracledb.defaultConnection();
    // preserve old module and action
    let result = conn.execute(
        `BEGIN
dbms application info.read module(:1 prev module, :1 prev action);
            l prev module: {
                dir: oracledb.BIND OUT,
                type: oracledb.STRING
            },
            l prev action: {
                dir: oracledb.BIND OUT,
                type: oracledb.STRING
            }
    );
    const prev module = result.outBinds.l prev module;
    const prev action = result.outBinds.l prev action;
    // set new module and action for the purpose of instrumenting
this
    // function.
    result = conn.execute(
            dbms application info.set module('plsql js mod',
'plsqlExample');
         END; `
    );
    // simulate some intensive processing ... While this is ongoing
    // and action in v$session should have changed to the values set
earlier
```



```
result = conn.execute(
         `BEGIN
            dbms session.sleep(60);
         END; `
    );
    // and finally reset module and action to what they were before
    result = conn.execute(
        `BEGIN
dbms application info.set module(:1 prev module, :1 prev action);
         END; `,
            l prev module: {
                dir: oracledb.BIND IN,
                type: oracledb.STRING,
                val: prev module
            },
            l prev action: {
                dir: oracledb.BIND IN,
                type: oracledb.STRING,
                val: prev action
            }
    );
}
```

This code demonstrates the use of OUT variables in the first call to DBMS_APPLICATION_INFO. The values can be retrieved using the result.outBinds object.

After storing the values in local variables, additional anonymous PL/SQL blocks are invoked, first setting module and action before entering a 60 second sleep cycle simulating processing of data. Once the simulated data processing routine finishes, the module and action are reset to their original values using named ${\tt IN}$ bind variables. Using bind variables is more secure than string concatenation.

Setting the module and action is a great way of informing the database about ongoing activity and allows for better grouping of activity in performance reports.

6.6 Error Handling in SQL Statements

JavaScript provides an exception framework like Java. Rather than returning an Error object as a promise or callback as in node-oracledb, the MLE JavaScript driver resorts to throwing errors. This concept is very familiar to PL/SQL developers.

Using try-catch-finally in JavaScript code is similar to the way PL/SQL developers use beginexception-end blocks to trap errors during processing.

Use the JavaScript throw() command if an exception should be re-thrown. This causes the error to bubble-up the stack after it has been dealt with in the catch block. Example 6-15 demonstrates this concept.

Example 6-14 SQL Error Handling Inside a JavaScript Function

CREATE OR REPLACE MLE MODULE js_err_handle_mod LANGUAGE JAVASCRIPT AS import oracledb from "mle-js-oracledb";

Create a procedure, errorHandling, using the module js err handle mod as follows:

```
CREATE OR REPLACE PROCEDURE errorHandling
AS MLE MODULE js_err_handle_mod
SIGNATURE 'errorHandling()';
/
```

When the procedure is executed, the console output indicates that the insert statement throws an error:

```
BEGIN
    errorHandling;
END;
/
An error occurred while inserting into thisTableDoesNotExist
```

In this case, the error is caught and processing stopped. No visible errors are shown to the consumer of the procedure, leading to a better user experience.

Example 6-15 Error Handling Using JavaScript throw() Command

This example demonstrates the use of the JavaScript throw() command in the catch block. Unlike the screen output shown for <code>js_err_handle_mod</code> in Example 6-14, a calling PL/SQL block will have to catch the error and either treat it accordingly or raise it again.



6.7 Working with JSON Data

The use of JSON data as part of a relational structure, more specifically the use of JSON columns in (relational) tables, is described.

Oracle Database supports JSON natively with relational database features, including transactions, indexing, declarative querying, and views. Unlike relational data, JSON data can be stored in the database, indexed, and queried without any need for a schema.¹

Oracle also provides a family of Simple Oracle Document Access (SODA) APIs for access to JSON data stored in the database. SODA is designed for schemaless application development without knowledge of relational database features or languages such as SQL and PL/SQL. It lets you create and store collections of documents in Oracle Database, retrieve them, and query them without needing to know how the documents are stored in the database.

JSON data is widely used for exchanging information between the application tier and database. Oracle REST Data Services (ORDS) is the most convenient tool for making REST calls to the database. Example 6-16 demonstrates this concept.

Manipulating JSON is one of JavaScript's core capabilities. Incoming JSON documents don't require parsing using <code>JSON.parse()</code>, they can be used straight away. Micro-service architectures greatly benefit from the enhanced options offered by JavaScript in the database.

See Also:

- Working with SODA Collections in MLE JavaScript Code for a detailed discussion of SODA and JavaScript in the database
- Oracle Database JSON Developer's Guide for information about the use of JSON in Oracle Database

Example 6-16 Inserting JSON Data into a Database Table

This example assumes that a REST API has been published in ORDS, allowing users to POST JSON data to the database. This way, administrators have an option to upload further

A JSON schema is not to be confused with the concept of a database schema: a database schema in Oracle Database is a separate namespace for database users to create objects such as tables, indexes, views, and many others without risking naming collisions.



departments into the <code>departments</code> table. Once the JSON data has been received, the MLE module uses ${\tt JSON_TABLE}$ () to convert the JSON data structure into a relational model.

```
CREATE OR REPLACE MLE MODULE RestApiModule LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function insertDepartments(depts) {
    // the input parameter must be valid JSON. This is ensured by using
    // JSON as a parameter in the call specification
    const conn = oracledb.defaultConnection();
    // convert JSON input to relational data and insert into a table
    const result = conn.execute(`
        INSERT INTO departments (
            department_id,
            department name,
            manager id,
            location id
        )
        SELECT
            jt.*
        FROM json table(:depts, '$[*]' columns
            department id path '$.department id',
            department name path '$.department name',
            manager_id path '$.manager_id',
location_id path '$.location_id'
        ) jt`,
            depts: {
                val: depts,
                 type: oracledb.DB TYPE JSON
        }
    );
}
```

Assuming the following input document is used for depts in the insertDepartments function, four rows will be inserted into the departments table:

```
[
    "department_id": 1010,
    "department_name": "New Department 1010",
    "manager_id": 200,
    "location_id": 1700
},
{
    "department_id": 1020,
    "department_name": "New Department 1020",
```

```
"manager_id": 201,
    "location_id": 1800
},
{
    "department_id": 1030,
    "department_name": "New Department 1030",
    "manager_id": 114,
    "location_id": 1700
},
{
    "department_id": 1040,
    "department_name": "New Department 1040",
    "manager_id": 203,
    "location_id": 2400
}
```

6.8 Using Large Objects (LOB) with MLE

A PL/SQL wrapper type is used to handle CLOBs and BLOBs with the MLE JavaScript driver. Handling of large objects such as CLOBs (Character Large Object) and BLOBs (Binary Large Object) with the MLE JavaScript driver differs from the node-oracledb driver. Rather than using a Node.js Stream interface, a PL/SQL wrapper type is used instead. The wrapper types for BLOBs and CLOBs are called OracleBlob and OracleClob respectively. Their use requires importing the objects from mle-js-plsqltypes.



BFILE, commonly counted among LOBs, is not supported.

See Also:

Server-Side JavaScript API Documentation for more information about mle-js-plsqltypes and the other JavaScript built-in modules

Topics

Writing LOBs

An example is used to show how to initialize and then write to a CLOB that is finally inserted into a table.

Reading LOBs

An example is used to show how to select a CLOB and then use the fetchInfo property to read the contents of the CLOB as a string.

6.8.1 Writing LOBs

An example is used to show how to initialize and then write to a CLOB that is finally inserted into a table.

Example 6-17 Inserting a CLOB into a Table

This example demonstrates how to insert a CLOB into a table. The table defines two columns: an ID column to be used as a primary key and a CLOB column named "C".

```
CREATE TABLE mle lob example (
    id NUMBER GENERATED ALWAYS AS IDENTITY,
    CONSTRAINT pk mle blob table PRIMARY KEY(id),
    c CLOB
);
CREATE OR REPLACE MLE MODULE ins clob mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
import { OracleClob } from "mle-js-plsqltypes";
export function clobInsert() {
    const conn = oracledb.defaultConnection();
    let mleCLOB = OracleClob.createTemporary(false);
   mleCLOB.open(OracleClob.LOB READWRITE);
   mleCLOB.write(
        1.
        'This is a CLOB and it has been inserted by the MLE JavaScript
SQL Driver'
    );
    const result = conn.execute(
        `INSERT INTO mle lob example (c) VALUES (:mleCLOB)`,
            mleCLOB: {
                type: oracledb.ORACLE CLOB,
                dir: oracledb.BIND IN,
                val: mleCLOB
            }
        }
    );
    conn.commit();
   mleCLOB.close();
}
```

CLOBs and BLOBs are defined in mle-js-plsqltypes. Rather than importing the entire module's namespace into ins_clob_mod only the relevant CLOB interface has been imported.

Therefore, the first step after grabbing a connection handle is to create a temporary uncached LOB locator. Following the successful initialization of the LOB, it is opened for read and write operations. A string is written to the CLOB with an offset of 1. Until this point, the LOB exists in memory. The call to connection.execute() inserts the CLOB in the table. Finally the transaction is committed and the LOB handle closed.



6.8.2 Reading LOBs

An example is used to show how to select a CLOB and then use the fetchInfo property to read the contents of the CLOB as a string.

Reading an LOB from the database is no different from reading other columns. Example 6-18 demonstrates how to fetch the row inserted by module <code>ins_clob_mod</code>, defined in Example 6-17.

Example 6-18 Read an LOB

```
CREATE OR REPLACE MLE MODULE read lob mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function clobSelect(id) {
    const conn = oracledb.defaultConnection();
    const result = conn.execute(
        `SELECT c FROM mle lob example WHERE id = :id`,
            id: {
                type: oracledb.NUMBER,
                dir: oracledb.BIND IN,
                val: id
        },
            fetchInfo: {
                "C": { type: oracledb.STRING }
            outFormat: oracledb.OBJECT
    );
    if (result.rows.length === 0) {
        throw Error(`clobSelect(): no data found for ID ${id}`);
        for (let row of result.rows) {
            console.log(
                 `The following text has been retrieved for ID ${id}:
                 ${row.C}`
            );
        }
    }
}
```

The only function, <code>clobSelect</code>, in <code>read_lob_mod</code> receives an ID as a parameter. It is used in the select statement's <code>where clause</code> as a bind variable to identify a row containing the CLOB. The <code>fetchInfo</code> property passed using <code>connection.execute()</code> instructs the database to fetch the CLOB as a string.

6.9 API Differences Between node-oracledb and mle-js-oracledb

There are several differences between node-oracledb and mle-js-oracledb, including the methods for handling connection management and type mapping.

See Also:

Server-Side JavaScript API Documentation for more information about JavaScript built-in modules

Topics

- Synchronous API and Error Handling
- Connection Handling
- Transaction Management
- Type Mapping
- Unsupported Data Types
- Miscellaneous Features Not Available with the MLE JavaScript SQL Driver

6.9.1 Synchronous API and Error Handling

Compared to node-oracledb, the mle-js-oracledb driver operates in a synchronous mode, throwing exceptions as they happen. If an asynchronous behavior is desired, calls to mle-js-oracledb can be wrapped into async functions.

During synchronous operations, API calls block until either a result or an error are returned. Errors caused by SQL execution are reported as JavaScript exceptions, otherwise they return the same properties as the node-oracledb Error object.

The following methods neither return a promise nor do they take a callback parameter. They either return the result or throw an exception.

- connection.execute
- connection.executeMany
- connection.getStatementInfo
- connection.getSodaDatabase
- connection.commit
- connection.rollback
- resultset.close
- resultset.getRow
- resultset.getRows

The following method cannot be implemented in a synchronous way and is omitted in the MLE JavaScript driver.

connection.break

node-oracledb provides a LOB (Large Object) class to provide streaming access to LOB types. The LOB class implements the asynchronous Node.js Stream API and cannot be supported in the synchronous MLE JavaScript environment. Large objects are supported using an alternative API in the MLE JavaScript driver. For these reasons, the following LOB-related functionality is not supported.



- connection.createLob
- property oracledb.lobPrefetchSize
- constant oracledb.BLOB
- constant oracledb.CLOB

node-oracledb also implements asynchronous streaming of query results, another feature that's based on the Node.js Stream API. A streaming API cannot be represented in a synchronous interface as used by the MLE JavaScript driver, therefore the following functionality is not available.

- connection.queryStream()
- resultSet.toQueryStream()

6.9.2 Connection Handling

The method of connection handling with the MLE JavaScript driver is described. All SQL statements that are executed via the server-side MLE JavaScript driver are executed in the current session that is running the JavaScript program. SQL statements are executed with the privileges of the user on whose behalf JavaScript code is executed. As in the node-oracledb API, JavaScript code using the MLE JavaScript driver must acquire a Connection object to execute SQL statements. However, the only connection available is the implicit connection to the current database session.

JavaScript code must acquire a connection to the current session using the MLE-specific oracledb.defaultConnection() method. On each invocation, it returns a connection object that represents the session connection. Creation of connections with the oracledb.createConnection method of node-oracledb is not supported by the MLE JavaScript driver; neither is the creation of a connection pool supported. Connection objects are implicitly closed and so the call to connection.close() is not available with the MLE JavaScript driver.

There is also no statement cursor caching with the MLE JavaScript driver and therefore there is no stmtCacheSize property.

The Real Application Cluster (RAC) option offers additional features, designed to increase availability of applications. These include Fast Application Notification (FAN) and Runtime Load Balancing (RLB), neither of which are supported by the MLE JavaScript driver.

6.9.3 Transaction Management

With respect to transaction management, server-side MLE JavaScript code behaves exactly like PL/SQL procedures and functions.

A JavaScript program is executed in the current transaction context of the calling SQL or PL/SQL statement. An ongoing transaction can be controlled by executing COMMIT, SAVEPOINT, or ROLLBACK commands. Alternatively, the methods connection.commit() and connection.rollback() can be used.

MLE JavaScript SQL driver connections cannot be explicitly closed. Applications relying on node-oracledb behavior where closing a connection performs a rollback of the transaction will need adjusting. The MLE JavaScript SQL driver neither performs implicit commit nor rollback of transactions.



The node-oracledb driver features an auto-commit flag, defaulting to false. The MLE JavaScript SQL driver does not implement this feature. If specified, the connection.execute() function ignores the parameter.

6.9.4 Type Mapping

The MLE JavaScript driver adheres to the behavior of node-oracledb with respect to conversions between PL/SQL types and JavaScript types.

By default, PL/SQL types map to native JavaScript types (except for BLOBs and CLOBs). Values fetched from query results are implicitly converted. See MLE Type Conversions for more details about MLE type mappings.

As with node-oracledb, the conversion from non-character data types and vice versa is directly impacted by the NLS session parameters. The MLE runtime locale has no impact on these conversions.

To avoid loss of precision when converting between native JavaScript types and PL/SQL data types, the MLE JavaScript driver introduces new wrapper types.

- oracledb.ORACLE NUMBER
- oracledb.ORACLE CLOB
- oracledb.ORACLE BLOB
- oracledb.ORACLE TIMESTAMP
- oracledb.ORACLE TIMESTAMP TZ
- oracledb.ORACLE DATE
- oracledb.ORACLE INTERVAL YM
- oracledb.ORACLE INTERVAL DS

As with node-oracledb, the default mapping to JavaScript types may be overridden on a case-by-case basis using the fetchInfo property on connection.execute(). Type constants like oracledb.ORACLE_NUMBER may be used to override the type mapping for a specific NUMBER column in order to avoid implicit conversion and loss of precision.

Additionally, the JavaScript MLE SQL driver provides a way to change the default mapping of PL/SQL types globally. If the oracledb.fetchAsPlsqlWrapper property contains the corresponding type constant, Oracle values are fetched as SQL wrapper types previously described. As with the existing property oracledb.fetchAsString, this behavior can be overridden using fetchInfo and oracledb.DEFAULT. Because MLE JavaScript does not support a Buffer class, and instead uses Uint8Array, property oracledb.fetchAsBuffer from node-oracledb does not exist in mle-js-oracledb, which instead uses oracledb.fetchAsUint8Array.

Changing the type mapping to fetch JavaScript SQL wrapper types by default accounts for the following scenarios:

- Oracle values are mainly moved between queries and DML statements, so that the type conversions between PL/SQL and JavaScript types are an unnecessary overhead.
- It is crucial to avoid data loss.



Example 6-19 Using JavaScript Native Data Types vs Using Wrapper Types

This example demonstrates the effect of using JavaScript native data types for calculations. It also compares the loss of precision using JavaScript native types versus using wrapper types.

```
CREATE OR REPLACE MLE MODULE js v wrapper mod LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function precisionLoss() {
    const conn = oracledb.defaultConnection();
    let result = conn.execute(`SELECT 0.1 summand1`);
   let summand1 = result.rows[0][0];
   result = conn.execute(`SELECT 0.2 summand2`);
    let summand2 = result.rows[0][0];
    console.log(
         Summand1 ${summand1} + summand2 ${summand2} = ${summand1 +
summand2}
    );
export function precisionPreservedGlobal() {
    const conn = oracledb.defaultConnection();
    oracledb.fetchAsPlsqlWrapper = [oracledb.NUMBER];
    let result = conn.execute(`SELECT 0.1 summand1`);
   let summand1 = result.rows[0][0];
    result = conn.execute(`SELECT 0.2 summand2`);
    let summand2 = result.rows[0][0];
    console.log(
        `Summand1 ${summand1} + summand2 ${summand2} = $
{summand1.add(summand2)}`
    );
}
```

When executing the above functions, the difference in precision becomes immediately obvious.

Rather than setting the global oracledb.fetchAsPlsqlWrapper property, it is possible to override the setting per invocation of connection.execute(). Example 6-20 shows how precisionPreservedGlobal() can be rewritten by setting precision inline.

Example 6-20 Overriding the Global oracledb.fetchAsPlsqlWrapper Property

This example extends Example 6-19 by showing how precisionPreservedGlobal() can be rewritten by preserving precision inline. It demonstrates that rather than setting the global oracledb.fetchAsPlsqlWrapper property, it is possible to override the setting per invocation of connection.execute().

```
CREATE OR REPLACE MLE MODULE ovr_glob_prop_mod LANGUAGE JAVASCRIPT AS import oracledb from "mle-js-oracledb";
```



```
export function precisionPreservedInline() {
    const conn = oracledb.defaultConnection();
    let result = conn.execute(
        `SELECT 0.1 summand1`,
        { fetchInfo: { "SUMMAND1": { type:
oracledb.ORACLE NUMBER } }
    );
    let summand1 = result.rows[0][0];
    result = conn.execute(
        `SELECT 0.2 summand2`,
        { fetchInfo: { "SUMMAND2": { type:
oracledb.ORACLE NUMBER } }
    );
    let summand2 = result.rows[0][0];
    console.log(
        `Summand1 ${summand1} + summand2 ${summand2} = $
{summand1.add(summand2)}`
```

6.9.5 Unsupported Data Types

The MLE JavaScript driver does not currently support these data types:

- LONG
- LONG RAW
- XMLType
- BFILE
- REF CURSOR

6.9.6 Miscellaneous Features Not Available with the MLE JavaScript SQL Driver

Differences between what features are available with the MLE JavaScript driver and with node-oracledb are described.

Error handling in the MLE JavaScript driver relies on the JavaScript exception framework rather than using a callback/promise as node-oracledb does. The error thrown by the MLE JavaScript SQL driver is identical to the Error object available with node-oracledb.

Several additional client-side features available in node-oracledb are not supported by the server-side MLE environment. The MLE JavaScript driver omits the API for these features.

The following features are currently unavailable:

Continuous Query Notification (CQN)



- Advanced Queuing is not supported natively, the PL/SQL API can be used as a workaround
- Connection.subscribe()
- Connection.unsubscribe()
- All Continuous Query Notification constants in the oracledb class
- All Subscription constants in the oracledb class



7

Working with SODA Collections in MLE JavaScript Code

Simple Oracle Document Access (SODA) is a set of NoSQL-style APIs that let you create and store collections of documents (in particular JSON) in Oracle Database, retrieve them, and query them, without needing to know Structured Query Language (SQL) or how the documents are stored in the database.

SODA APIs exist for different programming languages and include support for MLE JavaScript. SODA APIs are *document-centric*. You can use any SODA implementation to perform create, read, update, and delete (CRUD) operations on documents of nearly any kind (including video, image, sound, and other binary content). You can also use any SODA implementation to query the content of JavaScript Object Notation (JSON) documents using pattern-matching: query-by-example (QBE). CRUD operations can be driven by document keys or by QBEs.

This chapter covers JavaScript in the database, based on Multilingual Engine (MLE) as opposed to the client-side <code>node-oracledb</code> driver. Whenever JavaScript is mentioned in this chapter it implicitly refers to MLE JavaScript.



Oracle Database Introduction to Simple Oracle Document Access (SODA) for a complete introduction to SODA

Topics

- High-Level Introduction to Working with SODA for In-Database JavaScript
 The SODA API is part of the MLE JavaScript SQL driver. Interaction with collections and documents requires you to establish a connection with the database first, before a SODA database object can be obtained.
- SODA Objects
 Objects used with the SODA API.
- Using SODA for In-Database JavaScript
 How to access SODA for In-Database JavaScript is described, as well as how to use it to perform create, read (retrieve), update, and delete (CRUD) operations on collections.

7.1 High-Level Introduction to Working with SODA for In-Database JavaScript

The SODA API is part of the MLE JavaScript SQL driver. Interaction with collections and documents requires you to establish a connection with the database first, before a SODA database object can be obtained.

The SODA database is the top-level abstraction object when working with the SODA API.

Figure 7-1 demonstrates the standard control flow.

Figure 7-1 SODA for In-Database JavaScript Basic Workflow



Applications that aren't ported from client-side Node.js or Deno can benefit from coding aids available in the MLE JavaScript SQL driver, such as a number of frequently used variables that are available in the global scope. For a complete list of available global variables and types, see Working with the MLE JavaScript Driver.

For SODA applications the most important global variable is the <code>soda</code> object, which represents the <code>SodaDatabase</code> object. The availability of the <code>soda</code> object in the global scope reduces the need for writing boilerplate code.

7.2 SODA Objects

Objects used with the SODA API.

The following objects are at the core of the SODA API:

SodaDatabase: The top-level object for SODA operations. This is acquired from an Oracle Database connection. A SODA database is an abstraction, allowing access to SODA collections in that SODA database, which then allow access to documents in those collections. A SODA database is analogous to an Oracle

Database user or schema. A collection is analogous to a table. A document is analogous to a table row with one column for a unique document key, a column for the document content, and other columns for various document attributes. With the MLE JavaScript SQL driver, the soda object is available as a global variable, which represents the SodaDatabase object and reduces the need for writing boilerplate code.

- SodaCollection: Represents a collection of SODA documents. By default, collections allow JSON documents to be stored, and they add a default set of metadata to each document. This is recommended for most users. However optional metadata can set various details about a collection, such as its database storage, whether it should track version and time stamp document components, how such components are generated, and what document types are supported.
- SodaDocument: Represents a document. Typically, the document content will be JSON.
 The document has properties including the content, a key, timestamps, and the media type. By default, document keys are automatically generated.

When working with collections and documents stored therein, you will make use of the following objects:

- SodaDocumentCursor: A cursor object representing the result of the getCursor() method from a find() operation. It can be iterated over to access each SodaDocument.
- **SodaOperation:** An internal object used with find() to perform read and write operations on documents. Chained methods set properties on a SodaOperation object which is then used by a terminal method to find, count, replace, or remove documents. This is an internal object that should not be directly accessed.



Server-Side JavaScript API Documentation for information about using SODA objects with ${\tt mle-js-oracledb}$

7.3 Using SODA for In-Database JavaScript

How to access SODA for In-Database JavaScript is described, as well as how to use it to perform create, read (retrieve), update, and delete (CRUD) operations on collections.

This section describes SODA for MLE JavaScript. Code snippets in this section are sometimes abridged for readability. Care has been taken to ensure that JavaScript functions are listed in their entirety, but they aren't runnable on their own. Embedding the function definition into a JavaScript module and importing the MLE JavaScript SQL driver will convert these code examples to valid JavaScript code for Oracle Database 23c.

Topics

- Getting Started with SODA for In-Database JavaScript
 How to access SODA for In-Database JavaScript is described, as well as how to use it to create a database collection, insert a document into a collection, and retrieve a document from a collection.
- Creating a Document Collection with SODA for In-Database JavaScript
 How to use SODA for In-Database JavaScript to create a new document collection is explained.



- Opening an Existing Document Collection with SODA for In-Database JavaScript You can use the method SodaDatabase.openCollection() to open an existing document collection or to test whether a given name names an existing collection.
- Checking Whether a Given Collection Exists with SODA for In-Database JavaScript
 - You can use <code>SodaDatabase.openCollection()</code> to check for the existence of a given collection. It returns <code>null</code> if the collection argument does not name an existing collection; otherwise, it opens the collection having that name.
- Discovering Existing Collections with SODA for In-Database JavaScript You can use SodaDatabase.getCollectionNames() to fetch the names of all existing collections for a given SodaDatabase object.
- Dropping a Document Collection with SODA for In-Database JavaScript You use SodaCollection.drop() to drop an existing collection.
- Creating Documents with SODA for In-Database JavaScript
 Creation of documents by SODA for In-Database JavaScript is described.
- Inserting Documents into Collections with SODA for In-Database JavaScript SodaCollection.insertOne() or a related call such as sodaCollection.insertOneAndGet() offers convenient ways to add documents to a collection. These methods create document keys automatically, unless the collection is configured with client-assigned keys and the input document provides the key.
- Saving Documents into Collections with SODA for In-Database JavaScript
 You use SodaCollection.save() and saveAndGet() to save documents into
 collections.
- SODA for In-Database JavaScript Read and Write Operations
 The primary way you specify read and write operations (other than insert and save) is to use methods provided by the SodaOperation class. You can chain together SodaOperation methods to specify read or write operations against a collection.
- Finding Documents in Collections with SODA for In-Database JavaScript
 To find documents in a collection, you invoke SodaCollection.find(). It creates
 and returns a SodaOperation object which is used via method chaining with
 nonterminal and terminal methods.
- Replacing Documents in a Collection with SODA for In-Database JavaScript
 To replace the content of one document in a collection with the content of another,
 you start by looking up the document to be modified using its key. Because
 SodaOperation.key() is a nonterminal operation, the easiest way to replace the
 contents is to chain SodaOperation.key() to SodaOperation.replaceOne() or
 SodaOperation.replaceOneAndGet().
- Removing Documents from a Collection with SODA for In-Database JavaScript Removing documents from a collection is similar to replacing. The first step is to perform a lookup operation, usually based on the document's key or by using a search expression in SodaOperation.filter(). The call to SodaOperation.remove() is a terminal operation, in other words the last operation in the chain.
- Indexing the Documents in a Collection with SODA for In-Database JavaScript Indexes can speed up data access, regardless of whether you use the NoSQL style SODA API or a relational approach. You index documents in a SODA



collection using SodaCollection.createIndex(). Its IndexSpec parameter is a textual JSON index specification.

- Getting a Data Guide for a Collection with SODA for In-Database JavaScript
 A data guide is a summary of the structural and type information contained in a set of
 JSON documents. It records metadata about the fields used in those documents. They
 provide great insights into JSON documents and are invaluable for getting an overview of
 a data set.
- Handling Transactions with SODA for In-Database JavaScript
 Unlike the client-side JavaScript SQL driver, the MLE JavaScript SQL driver does not
 provide an autoCommit feature. You need to commit or roll your transactions back, either
 in the PL/SQL layer in case of module calls, or directly in the JavaScript code by calling
 connection.commit() or connection.rollback().
- Summary Example of Using SODA for In-Database JavaScript
 An end-to-end example is included to demonstrate using SODA for In-Database JavaScript from start to finish.

7.3.1 Getting Started with SODA for In-Database JavaScript

How to access SODA for In-Database JavaScript is described, as well as how to use it to create a database collection, insert a document into a collection, and retrieve a document from a collection.

Before you can get started working with SODA for MLE JavaScript, the account used for storing collections must be granted the SODA APP roles:

```
grant soda app to soda user
```

Accessing SODA functionality requires the use of the MLE JavaScript SQL driver. Because the database session exists by the time the code is invoked, no additional connection handling is necessary. Example 7-1 demonstrates how to:

- Create a SODA collection,
- Insert a JSON document into it, and
- Iterate over all SODA Documents in the collection, printing their contents on screen

Each concept presented by Example 7-1 - creating collections, adding and modifying documents, and dropping collections - is addressed in more detail later in this chapter.

Example 7-1 SODA with MLE JavaScript General Workflow

This example demonstrates the general workflow using SODA collections with MLE JavaScript.

```
CREATE OR REPLACE MLE MODULE intro_soda_mod
LANGUAGE JAVASCRIPT AS

import oracledb from "mle-js-oracledb";

export function testSODA(dropCollection) {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();
```



```
// get a SODA database
   const db = connection.getSodaDatabase();
   // Create a collection with the name "MyJSONCollection".
   // This creates a database table, also named "MyJSONCollection",
   // to store the collection. If a collection with the same name
   // exists it will be opened
   const col = db.createCollection("MyJSONCollection");
   // Create a JSON document (based on the HR.EMPLOYEES table for
employee 100)
   const doc = {
       "employee id": 100,
       "job id": "AD PRES",
       "last name": "King",
       "first name": "Steven",
       "email": "SKING",
       "manager id": null,
       "department id": 90
   };
   // Insert the document into a collection.
   col.insertOne(doc);
    // Find all documents in the collection and print them on screen
   try {
       // use a cursor to iterate over all documents in the collection
       const c = col.find().getCursor();
       let resultDoc;
       while (resultDoc = c.getNext()) {
           const content = resultDoc.getContent();
           console.log(`
               _____
                              ${resultDoc.key}
               content (select fields):
               - employee_id ${content.employee_id}
               - job_id ${content.job_id}
- name ${content.first.name}
                - name
                            ${content.first name} $
{content.last name}
               version: ${resultDoc.version}
               last modified: ${resultDoc.lastModified}
               created on: ${resultDoc.createdOn}
               media type: ${resultDoc.mediaType}`
           );
       }
       // it is very important to close the SODADocumentCursor to
free resources
       c.close();
   } catch (err) {
       throw err;
   }
```



```
// commit open transactions
connection.commit();

// optionally drop the collection
if (dropCollection) {
    col.drop();
}
```

7.3.2 Creating a Document Collection with SODA for In-Database JavaScript

How to use SODA for In-Database JavaScript to create a new document collection is explained.

Collections allow you to logically group documents. Before a collection can be created or accessed, a few more steps must be completed. Begin by creating a connection object. The connection object is the starting point for all SODA interactions in the MLE JavaScript module:

```
// get a connection handle to the database session
const connection = oracledb.defaultConnection();
```

Once the connection is obtained, you can use it to call <code>Connection.getSodaDatabase()</code>, a prerequisite for creating the collection:

```
// get a SODA database
const db = connection.getSodaDatabase();
```

With the SODA database available, the final step is to create the collection. Note that collection names are case-sensitive:

```
// Create a collection with the name "MyJSONCollection".
// This creates a database table, also named "MyJSONCollection",
// to store the collection. If a collection with the same name
// exists, it will be opened
const col = db.createCollection("MyJSONCollection");
```

The preceding statement creates a collection that, by default, allows JSON documents to be stored. If the collection name passed to <code>SodaDatabase.createCollection()</code> is that of an existing collection, it will simply be opened. You can alternatively open a known, existing collection using <code>SodaDatabase.openCollection()</code>.

Unless custom metadata (as seen in Example 7-2) is provided to SodaDatabase.createCollection(), default collection metadata will be supplied. The default metadata has the following characteristics:

- Each document in the collection has these components:
 - Key



- Content
- Creation timestamp
- Last-modified timestamp
- Version
- The collection can store only JSON documents.
- Document keys and version information are generated automatically using a Universally Unique Identifier (UUID) function.

Optional collection metadata can be provided to the call to <code>createCollection()</code>, however, the default collection configuration is recommended in most cases.

If a collection with the same name already exists, it is simply opened and its object is returned. If custom metadata is passed to the method and does not match that of the existing collection, the collection is not opened and an error is raised. To match, all metadata fields must have the same values.



Oracle Database Introduction to Simple Oracle Document Access (SODA) for more details about collection metadata.

Example 7-2 Creating a Collection with Custom Metadata

This example demonstrates how to create a collection with custom metadata. Rather than storing the key as an automatically created UUID value stored in a VARCHAR2 column, the key is stored as a numeric field.

```
export function createCustomCollection() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();
    // get a SODA database
    const db = connection.getSodaDatabase();
    // create a collection with custom metadata. Unless a custom
    // metadata field (keyColumn, versionColumn) is specified it
    // will not be used/present in the collection. The metadata
    // argument is mostly based on the defaults except for the
    // client-assigned key column
    const col = db.createCollection(
        "myCustomCollection",
            "metaData": {
                "keyColumn": {
                    "name": "ID",
                    "sqlType": "NUMBER",
                    "assignmentMethod": "CLIENT"
                },
                "versionColumn": {
                    "name": "VERSION",
```



Although the keyColumn is the only field whose value deviates from the defaults, it is mandatory to add versionColumn, lastModifiedColumn, and creationTimeColumn or they won't be used in the collection.

7.3.3 Opening an Existing Document Collection with SODA for In-Database JavaScript

You can use the method <code>SodaDatabase.openCollection()</code> to open an existing document collection or to test whether a given name names an existing collection.

Example 7-3 Opening an Existing Document Collection

This example opens the collection named <code>collectionName</code>. It is very important to check that the collection object returned by <code>SodaDatabase.openCollection()</code> is not <code>null</code>. Rather than throwing an error, the method will return a <code>null</code> value should the requested collection not exist.

```
export function openCollection(collectionName) {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();

    // perform a lookup. If a connection cannot be found by that
    // name no exception nor error are thrown, but the resulting
    // collection object will be null
    const col = db.openCollection(collectionName);
    if (col === null) {
        throw `No such collection ${collectionName}`;
    }

    // do something with the collection
}
```

7.3.4 Checking Whether a Given Collection Exists with SODA for In-Database JavaScript

You can use <code>SodaDatabase.openCollection()</code> to check for the existence of a given collection. It returns <code>null</code> if the collection argument does not name an existing collection; otherwise, it opens the collection having that name.

In Example 7-3, if collectionName does not name an existing collection then col is assigned the value null.

7.3.5 Discovering Existing Collections with SODA for In-Database JavaScript

You can use SodaDatabase.getCollectionNames() to fetch the names of all existing collections for a given SodaDatabase object.

If the number of collections is very large, you can limit the number of names returned. Additionally, the lookup can be limited to collections starting with a user-defined string as demonstrated by Example 7-5.

Example 7-4 Fetching All Existing Collection Names

This example prints the names of all existing collections using the method getCollectionNames().

```
export function printCollectionNames() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();

    // loop over all collection names
    const allCollections = db.getCollectionNames();
    for (const col of allCollections) {
        console.log(`- ${col}`);
    }
}
// defaultConnection();
```

Example 7-5 Filtering the List of Returned Collections

This example limits the results of getCollectionNames() by only printing the names of collections that begin with a user-defined string, startWith.

```
export function printSomeCollectionNames(numHits, startWith) {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();
    // get a SODA database
```



```
const db = connection.getSodaDatabase();
   // loop over all collection names
   const allCollections = db.getCollectionNames(
           limit: numHits,
            startsWith: startWith
   );
   for (const col of allCollections) {
       console.log(`- ${col}`);
}
```

7.3.6 Dropping a Document Collection with SODA for In-Database **JavaScript**

You use SodaCollection.drop() to drop an existing collection.



Caution:

Do not use SQL to drop the database table that underlies a collection. Dropping a collection involves more than just dropping its database table. In addition to the documents that are stored in its table, a collection has metadata, which is also persisted in Oracle Database. Dropping the table underlying a collection does not also drop the collection metadata.

Note:

Day-to-day use of a typical application that makes use of SODA does not require that you drop and re-create collections. But if you need to do that for any reason then this guideline applies.

Do not drop a collection and then re-create it with different metadata if there is any application running that uses the collection in any way. Shut down any such applications before re-creating the collection, so that all live SODA objects are released.

There is no problem just dropping a collection. Any read or write operation on a dropped collection raises an error. And there is no problem dropping a collection and then re-creating it with the same metadata. But if you re-create a collection with different metadata, and if there are any live applications using SODA objects, then there is a risk that a stale collection is accessed, and no error is raised in this case.

In SODA implementations that allow collection metadata caching, such as SODA for Java, this risk is increased if such caching is enabled. In that case, a (shared or local) cache can return an entry for a stale collection object even if the collection has been dropped.





Commit all writes to a collection before using <code>SodaCollection.drop()</code>. For the method to succeed, all uncommitted writes to the collection must first be committed. Otherwise, an exception is raised.

Example 7-6 Dropping a Collection

This example drops collection col.

```
CREATE OR REPLACE MLE MODULE drop_col_mod
LANGUAGE JAVASCRIPT AS

import oracledb from "mle-js-oracledb";

export function openAndDropCollection(collectionName) {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();

    // look the collection up
    const col = db.openCollection(collectionName);
    if (col === null) {
        throw `No such collection ${collectionName}`;
    }

    // drop the collection - POTENTIALLY DANGEROUS
    col.drop();
}
```

7.3.7 Creating Documents with SODA for In-Database JavaScript

Creation of documents by SODA for In-Database JavaScript is described.

The SodaDocument class represents SODA documents. Although its focus is on JSON documents, it supports other content types as well. A SodaDocument stores both the actual document's contents as well as metadata.

JavaScript is especially well-suited to work with JSON by design, giving it an edge over other programming languages.

Here is an example of a simple JSON document:

```
// Create a JSON document (based on the HR.EMPLOYEES table for
employee 100)
const doc = {
    "employee_id": 100,
    "job_id": "AD_PRES",
```



```
"last_name": "King",
  "first_name": "Steven",
  "email": "SKING",
  "manager_id": null,
  "department_id": 90
};
```



In SODA, JSON content must conform to RFC 4627.

SodaDocument objects can be created in three ways:

- As a result of sodaDatabase.createDocument(). This is a proto-SodaDocument object
 usable for SODA insert and replace methods. The SodaDocument will have content and
 media type components set. Attributes like createdOn will not be defined. Optional
 attributes not specified when calling SodaDocument.createDocument() will not be defined
 either.
- As a result of a read operation from the database, such as calling sodaOperation.getOne(), or from sodaDocumentCursor.getNext() after a sodaOperation.getCursor() call. These return complete SodaDocument objects containing the document content and attributes, such as time stamps.
- As a result of sodaCollection.insertOneAndGet(), sodaCollection.insertManyAndGet() methods. These return SodaDocuments that contain all attributes except the document content itself. They are useful for finding document attributes such as system generated keys, and versions of new and updated documents.

A document has these components:

- Kev
- Content
- Creation time stamp
- Last-modified time stamp
- Version
- Media type ("application/json" for JSON documents)

Example 7-7 Creating SODA Documents

```
export function createJSONDoc() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();

    // define a payload
    const payload = {
```



```
"employee id": 100,
        "job id": "AD PRES",
        "last name": "King",
       "first name": "Steven",
        "email": "SKING",
        "manager id": null,
       "department id": 90
    };
    // create a SODA document using the default key
    const docDefaultKey = db.createDocument(payload);
    console.log(`
       SODA Document using default key
       content (select fields):
       - employee id ${docDefaultKey.getContent().employee id}
       - job id ${docDefaultKey.getContent().job id}
       - first name ${docDefaultKey.getContent().first name}
       media type: ${docDefaultKey.mediaType}`
    );
    // create a SODA document using a custom key (requires that the
    // collection this document is inserted to uses a user-assigned
key)
    const docCustomKey = db.createDocument(
       payload,
       {
           key: "12345"
    );
    // print details about the new SODA Document
    console.log(`
       ______
       SODA Document using custom key
       content (select fields):
       - employee id ${docCustomKey.getContent().employee id}
       - job id ${docCustomKey.getContent().job id}
       - first name ${docCustomKey.getContent().first_name}
       key:
                    ${docCustomKey.key}
       media type: ${docCustomKey.mediaType}`
    );
}
```

Creating SodaDocument instances as shown in this example is the exception rather than the norm. In most cases, developers use SodaCollection.insertOne() or SodaCollection.insertOneAndGet(). The use of SodaCollection.insertOne() is demonstrated in Example 7-8.

7.3.8 Inserting Documents into Collections with SODA for In-Database JavaScript

SodaCollection.insertOne() or a related call such as sodaCollection.insertOneAndGet() offers convenient ways to add documents to a collection. These methods create document keys automatically, unless the collection is configured with client-assigned keys and the input document provides the key.

SodaCollection.insertOne() simply inserts the document into the collection, whereas SodaCollection.insertOneAndGet() additionally returns a result document. The resulting document contains the document key and any other generated document components, except for the actual document's content (this is done to improve performance).

Both methods automatically set the values of the creation time stamp, last-modified time stamp, and version, unless the collection has been created with custom metadata. Custom metadata might not include all the default metadata. When querying attributes not defined by the collection a null value is returned.



If the collection is configured with client-assigned document keys (which is not the default case), and the input document provides a key that identifies an existing document in the collection, then these methods throw an exception. If you want the input document to *replace* the existing document instead of causing an exception, see Saving Documents into Collections with SODA for In-Database JavaScript.

Example 7-8 Inserting a SODA Document into a Collection

This example demonstrates how to insert a document into a collection using SodaCollection.insertOne().

```
export function insertOneExample() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();

    // create/open a SODA collection
    const col = db.createCollection("MyJSONCollection");

    // define a payload
    const payload = {
        "employee_id": 100,
        "first_name": "Steven",
        "last_name": "King"
    };

    // insert the document
```



```
col.insertOne (payload);
}
```

Example 7-9 Inserting an Array of Documents into a Collection

This example demonstrates the use of SodaCollection.insertMany() to insert multiple documents with one command.

```
export function insertManyExample() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();
    // get a SODA database
    const db = connection.getSodaDatabase();
    // create/open a SODA collection
    const col = db.createCollection("MyJSONCollection");
    // define a payload (an array of objects)
    const payload = [
            "employee id": 100,
            "first name": "Steven",
            "last name": "King"
        },
            "employee id": 101,
            "first name": "Neena",
            "last name": "Kochhar"
        },
            "employee id": 102,
            "first name": "Lex",
            "last name": "De Haan"
        },
            "employee id": 103,
            "first name": "Alexander",
            "last name": "Hunold"
        },
            "employee id": 104,
            "first name": "Bruce",
            "last name": "Ernst"
        },
            "employee id": 105,
            "first name": "David",
            "last name": "Austin"
    ]
    // insert the document
```



```
col.insertMany(payload);
}
```

7.3.9 Saving Documents into Collections with SODA for In-Database JavaScript

You use SodaCollection.save() and saveAndGet() to save documents into collections.

These methods are similar to methods <code>insertOne()</code> and <code>insertOneAndGet()</code> except that, if the collection is configured with client-assigned document keys, and the input document provides a key that already identifies a document in the collection, then the input document <code>replaces</code> the existing document. In contrast, methods <code>insertOne()</code> and <code>insertOneAndGet()</code> throw an exception in that case.



By default, collections are configured with automatically generated document keys. Therefore, for a default collection, methods <code>save()</code> and <code>saveAndGet()</code> are equivalent to methods <code>insertOne()</code> and <code>insertOneAndGet()</code>, respectively.

Example 7-10 Saving a Document into a Collection

This example shows how to use <code>SodaCollection.saveAndGet()</code>. The JavaScript code stores a SODA document in a collection defined with custom metadata to ensure <code>save()</code> and <code>saveAndGet()</code> work as described. The key column specifically is user-managed, which is why you see the <code>SodaDocument</code> constructor mention a key.

```
CREATE OR REPLACE MLE MODULE save doc mod
LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function saveJSONDoc() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();
    // get a SODA database
    const db = connection.getSodaDatabase();
    // create/open a collection using custom metadata
    const col = db.openCollection("myCustomCollection");
    if ( col === null ) {
        throw `'myCustomCollection' does not exist`;
    // create the SODA document, use the employee id as its key
    let docCustomKey = db.createDocument(
        "employee_id": 100,
        "job id": "AD PRES",
```



```
"last name": "King",
    "first name": "Steven",
    "email": "SKING",
    "manager id": null,
    "department id": 90
    },
        "key": "100"
);
// insert the document into the custom collection
let savedDoc = col.saveAndGet(docCustomKey);
// get some meta information about the new document
console.log(`Document saved for the 1st time:
- department id: ${docCustomKey.getContent().department id}
- key:
                ${savedDoc.key}
- version:
                ${savedDoc.version}
- last modified: ${savedDoc.lastModified}
                ${savedDoc.createdOn}`);
- created on:
// correct the previous mistake - the document should have
// been inserted with a department id of 30
// - create the new document (important to use the same key)
docCustomKey = db.createDocument(
    "employee id": 100,
    "job id": "AD PRES",
    "last name": "King",
    "first name": "Steven",
    "email": "SKING",
    "manager id": null,
    "department id": 30
    },
        "key": "100"
);
// - update the existing document
savedDoc = col.saveAndGet(docCustomKey);
// - show the changed metadata. Cannot refer to the content
// as it is not returned by saveAndGet() for performance
// reasons
console.log(`Updated document saved successfully:
- department id: ${docCustomKey.getContent().department id}
- key:
                ${savedDoc.key}
- version:
                ${savedDoc.version}
- last modified: ${savedDoc.lastModified}
- created on: ${savedDoc.createdOn}`);
```

}

After the initial document has been saved to the collection, it turns out that the value for the document's department_id was wrong and needs correcting. The mistake is corrected and the document saved again, overwriting the previous contents and effectively setting the department id to the correct value of 30.

7.3.10 SODA for In-Database JavaScript Read and Write Operations

The primary way you specify read and write operations (other than insert and save) is to use methods provided by the <code>SodaOperation</code> class. You can chain together <code>SodaOperation</code> methods to specify read or write operations against a collection.

Nonterminal SodaOperation methods return the same object on which they are invoked, allowing them to be chained together.

A terminal SodaOperation method always appears at the end of a method chain to execute the operation.



A SodaOperation object is an internal object. You should not directly modify its properties.

Unless the <code>node-oracledb</code> documentation for a method says otherwise, you can chain together any nonterminal methods and you can end the chain with any terminal method. However, not all combinations make sense. For example, it does not make sense to chain method version() together with a method that does not uniquely identify the document, such as keys().

Table 7-1 Overview of Nonterminal Methods for Read Operations

Method	Description
key()	Find a document that has the specified document key.
keys()	Find documents that have the specified document keys.
filter()	Find documents that match a filter specification (a query-by-example expressed in JSON).
version()	Find documents that have the specified version. This is typically used with $key()$.
headerOnly()	Exclude document content from the result.
skip()	Skip the specified number of documents in the result.
limit()	Limit the number of documents in the result to the specified number.

Table 7-2 Overview of Terminal Methods for Read Operations

Method	Description
getOne()	Create and execute an operation that returns at most one document. For example, an operation that includes an invocation of nonterminal method key().



Table 7-2 (Cont.) Overview of Terminal Methods for Read Operations

Method	Description
getCursor()	Get a cursor over read operation results.
count()	Count the number of documents found by the operation.
getDocuments()	Gets an array of documents matching the query criteria.

Table 7-3 Overview of Terminal Methods for Write Operations

Method	Description
replaceOne()	Replace one document.
replaceOneAndGet()	Replace one document and return the result document.
remove()	Remove documents from a collection.

See Also:

- Node-oracledb Documentation for more details about the SodaOperations class.
- SODA Restrictions (Reference) for information about SODA restrictions.

7.3.11 Finding Documents in Collections with SODA for In-Database JavaScript

To find documents in a collection, you invoke <code>SodaCollection.find()</code>. It creates and returns a <code>SodaOperation</code> object which is used via method chaining with nonterminal and terminal methods.

To execute the query, obtain a cursor for its results by invoking <code>SodaOperation.getCursor()</code>. Then use the cursor to visit each document in the result list. This is illustrated by <code>Example 7-1</code> and other examples. It is important not to forget to close the cursor, to save resources.

However, this is not the typical workflow when searching for documents in a collection. It is more common to chain multiple methods provided by the SodaOperation class together.

Example 7-11 Finding a Document by Key

This example shows how to look up a document by its key using the methods find(), key(), and getOne().

There is an assumption that <code>customCollection</code> has been created as outlined in Example 7-2.

CREATE OR REPLACE MLE MODULE find_doc_mod LANGUAGE JAVASCRIPT AS



```
import oracledb from "mle-js-oracledb";
export function findDocByKey() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();
    // get a SODA database
    const db = connection.getSodaDatabase();
    // open a collection in preparation of a document
    // lookup
    const col = db.openCollection("myCustomCollection");
    if (col === null) {
        throw `myCustomCollection does not exist`;
    try {
        // perform a lookup of a document with key "100"
        const doc = col.find().key("100").getOne();
        console.log(`
            document found for key ${doc.key}
            contents: ${doc.getContentAsString()}`
        );
    } catch (err) {
        throw `No document found in 'myCustomCollection' with key 100`;
```

Note:

Keys need to be enclosed in double quotation marks even if they should be in numeric format.

In case the search for a given key fails, the database throws an ORA-01403 (no data found) exception. It is good practice to handle exceptions properly. In this example, the caller of the function has the responsibility to ensure the error is trapped and dealt with according to the industry's best-known methods.

Example 7-12 Looking up Documents Using Multiple Keys

This example uses the methods find(), keys(), getCursor(), and getNext() to search for multiple keys provided in an array.

```
export function findDocByKeys() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();
```



```
// open a collection in preparation of a document
    // lookup
   const col = db.openCollection("myCustomCollection");
    if (col === null) {
        throw `myCustomCollection does not exist`;
   try {
        // perform a lookup of a document with keys ["100","101"]
        const docCursor = col.find().keys(["100", "101"]).getCursor();
        let doc
        while (( doc = docCursor.getNext())) {
            console.log(`
                document found for key ${doc.key}
                contents: ${doc.getContentAsString()}`
            );
        }
       docCursor.close();
    } catch (err) {
        throw `No document found in 'myCustomCollection' with either
key`;
}
```

Rather than failing with an ORA-1403 (no data found exception), the find() operation simply doesn't return any data for a key not found in a collection. If none of the keys are found, nothing is returned.

Example 7-13 Using a QBE to Filter Documents in a Collection

This example uses filter() to locate documents in a collection. The nonterminal <code>SodaOperation.filter()</code> method provides a powerful way to filter JSON documents in a collection, allowing for complex document queries and ordering of JSON documents. Filter specifications can include comparisons, regular expressions, logical and spatial operators, among others.

The search expression defined in filterCondition matches all employees with an employee ID greater than 110 working in department 30.

```
export function findDocByFiltering() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();

    // open a collection in preparation of a document
    // lookup. This particular collection contains all the
    // rows from the HR.employees table converted to SODA
    // documents
    const col = db.openCollection("employees_collection");
    if (col === null) {
        throw `employees collection does not exist`;
```



```
}
   // find all employees with an employee id > 110 working
   // in department 30
   const filterCondition = {
       "employee id": { "$gt": 110 },
       "department id": "30"
   }
   try {
       // perform the lookup operation using the QBE
       const docCursor = col.find().filter(filterCondition).getCursor();
       while (( doc = docCursor.getNext())) {
           console.log(`
               document found matching the search criteria
                                ${doc.key}
               - employee id: ${doc.getContent().employee id}
               - department ID: ${doc.getContent().department id}
               - name:
                                ${doc.getContent().last name}`
           );
       }
       docCursor.close();
    } catch (err) {
       throw `No document found in 'myCustomCollection' with either key`;
}
```

See Also:

- Oracle Database Introduction to Simple Oracle Document Access (SODA) for an introduction to SODA filter specifications
- Oracle Database Introduction to Simple Oracle Document Access (SODA) for reference information about SODA filter specifications

Example 7-14 Using skip() and limit() in a Pagination Query

If the number of rows becomes too large, you may choose to paginate and or limit the number of documents returned. This example demonstrates using skip() and limit() in this type of circumstance.

```
export function paginationExample() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();
```



```
// open a collection in preparation of a document
    // lookup. This particular collection contains all the
    // rows from the HR.employees table converted to SODA
    // documents
    const col = db.openCollection("employees collection");
    if (col === null) {
        throw `employees collection does not exist`;
    }
    // find all employees with an employee id > 110 working
    // in department 80
    const filterCondition = {
        "employee id": { "$gt": 110 },
        "department id": "80"
    }
    try {
        // perform the lookup operation using the QBE, skipping the
first
        // 10 documents and limiting the result set to 10 documents
        const docCursor =
col.find().filter(filterCondition).skip(10).limit(10).getCursor();
        let doc
        while (( doc = docCursor.getNext())) {
           console.log(`
                document found matching the search criteria
                - key: ${doc.key}
                - employee id: ${doc.getContent().employee_id}`
           );
        }
        docCursor.close();
    } catch (err) {
        throw `No document found in 'myCustomCollection' with either
key`;
}
```

Example 7-15 Specifying Document Versions

This example uses the nonterminal version() method to specify a particular document version. This is useful for implementing optimistic locking, when used with the terminal methods for write operations.

```
export function versioningExample() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();
```



Example 7-16 Counting the Number of Documents Found

This example shows how to count the number of documents found in a collection using the find(), filter(), and count() methods. The filter() expression limits the result to all employees working in department 30.

```
export function countingExample() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();
    // get a SODA database
    const db = connection.getSodaDatabase();
    // open a collection in preparation of a document lookup
    const col = db.openCollection("employees collection");
    if (col === null) {
       throw `employees collection does not exist`;
    }
    try {
        // perform a lookup operation identifying all employees working
        // in department 30 limiting the result to headers only
        const filterCondition = {"department id": "30" };
        const numDocs = col.find().filter(filterCondition).count();
        console.log(`there are ${numDocs} documents matching the filter`);
    } catch (err) {
        throw `No document found in 'employees collection' matching the
filter`;
    }
```



7.3.12 Replacing Documents in a Collection with SODA for In-Database JavaScript

To replace the content of one document in a collection with the content of another, you start by looking up the document to be modified using its key. Because SodaOperation.key() is a nonterminal operation, the easiest way to replace the contents is to chain SodaOperation.key() to SodaOperation.replaceOne() or SodaOperation.replaceOneAndGet().

SodaOperation.replaceOne() merely replaces the document, whereas SodaOperation.replaceOneAndGet() replaces it and provides the resulting new document to the caller.

Both methods update the values of the last modified time stamp and the version. Replacement does not change the document key or the creation time stamp.

The difference between <code>SodaOperation.replace()</code> and <code>SodaOperation.save()</code> is that the latter performs an insert in case the key doesn't already exist in the collection. The replace operation requires an existing document to be found by the lookup via the <code>SodaOperation.key()</code> method.

Note:

Some version-generation methods, including the default method, generate hash values of the document content. In such a case, if the document content does not change then neither does the version.

Example 7-17 Replacing a Document in a Collection and Returning the Result Document

This example shows how to replace a document in a collection, returning a reference to the changed document.

```
CREATE OR REPLACE MLE MODULE rep_ret_doc_mod
LANGUAGE JAVASCRIPT AS

import oracledb from "mle-js-oracledb";

export function replaceExample() {

    // get a connection handle to the database session const connection = oracledb.defaultConnection();

    // get a SODA database const db = connection.getSodaDatabase();

    // open myCustomCollection const col = db.openCollection("myCustomCollection");
    if ( col === null ) {
        throw "'myCustomCollection' does not exist";
    }
}
```



```
// create the changed SODA document, change department id to 80
const newDoc = db.createDocument(
        "employee id": 100,
        "job id": "AD PRES",
        "last name": "King",
        "first name": "Steven",
        "email": "SKING",
        "manager id": null,
        "department id": 80
);;
// perform a lookup of the document about to be changed
    const resultDoc = col.find().key("100").replaceOneAndGet(newDoc);
    // print some metadata (note that content is not returned for
    // performance reasons)
    console.log(`Document updated successfully:
    - key:
                    ${resultDoc.key}
    - version:
                    ${resultDoc.version}
    - last modified: ${resultDoc.lastModified}
    - created on: ${resultDoc.createdOn}`);
} catch (err) {
    throw err;
```

Note:

Trying to read the changed contents will result in an error as the actual document's contents aren't returned, for performance reasons.

7.3.13 Removing Documents from a Collection with SODA for In-Database JavaScript

Removing documents from a collection is similar to replacing. The first step is to perform a lookup operation, usually based on the document's key or by using a search expression in SodaOperation.filter(). The call to SodaOperation.remove() is a terminal operation, in other words the last operation in the chain.

Example 7-18 Removing a Document from a Collection Using a Document Key

This example removes the document whose document key is "100".

```
CREATE OR REPLACE MLE MODULE rm_doc_mod LANGUAGE JAVASCRIPT AS import oracledb from "mle-js-oracledb";
```



```
export function removeByKey() {

    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();

    // open myCustomCollection
    const col = db.openCollection("myCustomCollection");
    if ( col === null ) {
        throw "'myCustomCollection' does not exist";
    }

    // perform a lookup of the document about to be removed and
    // ultimately remove it.
    const result = col.find().key("100").remove();
    if ( result.count === 0) {
        throw "Document was not deleted - is this an invalid key?"
    }
}
```

Example 7-19 Removing JSON Documents from a Collection Using a Filter

This example uses a filter to remove the JSON documents whose department_id is 70. It then prints the number of documents removed.

```
export function removeByFilter() {

    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();

    // open employees_collection
    const col = db.openCollection("employees_collection");
    if ( col === null ) {
        throw "'employees_collection' does not exist";
    }

    // perform a lookup based on a filter expression and remove the
    // documents matching the filter
    const result = col.find().filter({"department_id": "70"}).remove();
    return result.count;
}
```



7.3.14 Indexing the Documents in a Collection with SODA for In-Database JavaScript

Indexes can speed up data access, regardless of whether you use the NoSQL style SODA API or a relational approach. You index documents in a SODA collection using SodaCollection.createIndex(). Its IndexSpec parameter is a textual JSON index specification.

Existing indexes can be dropped using SodaCollection.dropIndex().

A JSON search index is used for full-text and ad hoc structural queries, and for persistent recording and automatic updating of JSON data-guide information.

See Also:

- Oracle Database Introduction to Simple Oracle Document Access (SODA) for an overview of using SODA indexing
- Oracle Database Introduction to Simple Oracle Document Access (SODA) for information about SODA index specifications
- Oracle Database JSON Developer's Guide for information about JSON search indexes
- Oracle Database JSON Developer's Guide for information about persistent data-guide information as part of a JSON search index

Example 7-20 Creating a B-Tree Index for a JSON Field with SODA for In-Database JavaScript

This example creates a B-tree non-unique index for numeric field department_id of the JSON documents in collection employees_collection.

```
CREATE OR REPLACE MLE MODULE b_tree_mod
LANGUAGE JAVASCRIPT AS

import oracledb from "mle-js-oracledb";

export function createBTreeIndex() {
    // get a connection handle to the database session const connection = oracledb.defaultConnection();

    // get a SODA database const db = connection.getSodaDatabase();

    // open the collection const col = db.openCollection("employees_collection");
    if ( col === null ) {
        throw "'employees_collection' does not exist";
    }

    const indexSpec = {
```



Example 7-21 Creating a JSON Search Index with SODA for In-Database JavaScript

This example shows how to create a JSON search index for indexing the documents in collection <code>employees_collection</code>. It can be used for ad hoc queries and full-text search (queries using QBE operator <code>\$contains</code>). It automatically accumulates and updates data-guide information about your JSON documents (aggregate structural and type information). The index specification has only field <code>name</code> (no *field* fields unlike the B-tree index in <code>Example 7-20</code>).

```
export function createSearchIndex() {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();
    // get a SODA database
    const db = connection.getSodaDatabase();
    // open the collection
    const col = db.openCollection("employees collection");
    if ( col === null ) {
        throw "'employees collection' does not exist";
    const indexSpec = {
        "name": "SEARCH AND DATA GUIDE IDX",
        "dataguide": "on",
        "search on": "text value"
    col.createIndex(indexSpec);
    console.log(`Search & Data Guide index successfully created`);
}
```

If you only wanted to speed up ad hoc (search) indexing, you should specify a value of "off" for field dataguide. The dataguide indexing feature can be turned off in the same way if it is not required.

Example 7-22 Dropping an Index with SODA for In-Database JavaScript

This example shows how you can drop an existing index on a collection using SodaCollection.dropIndex() and the force option.

```
export function dropIndex(indexName) {
    // get a connection handle to the database session
    const connection = oracledb.defaultConnection();

    // get a SODA database
    const db = connection.getSodaDatabase();

    // open the collection
    const col = db.openCollection("employees_collection");
    if ( col === null ) {
        throw "'employees_collection' does not exist";
    }

    // drop the index
    const result = col.dropIndex(indexName, {"force": true});
    if (! result.dropped ) {
        throw `Could not drop SODA index ${indexName}`;
    }
}
```

SodaCollection.dropIndex() returns a result object containing a single field: dropped. Its value is true if the index has been dropped, otherwise its value is false. The method succeeds either way.

An optional parameter object can be supplied to the method. Setting force to true forces dropping of a JSON index if the underlying Oracle Database domain index does not permit normal dropping.

7.3.15 Getting a Data Guide for a Collection with SODA for In-Database JavaScript

A data guide is a summary of the structural and type information contained in a set of JSON documents. It records metadata about the fields used in those documents. They provide great insights into JSON documents and are invaluable for getting an overview of a data set.

You can create a data guide using <code>SodaCollection.getDataGuide()</code>. To get a data guide in <code>SODA</code>, the collection must be <code>JSON-only</code> and have a <code>JSON</code> search index where the "dataguide" option is "on". Data guides are returned from <code>sodaCollection.getDataGuide()</code> as <code>JSON</code> content in a <code>SodaDocument</code>. The data guide is inferred from the collection as it currently is. As a collection grows and documents change, a new data guide is returned each subsequent time <code>getDataGuide()</code> is called.



Example 7-23 Generating a Data Guide for a Collection

This example gets a data guide for the collection <code>employees_collection</code> using the method <code>getDataGuide()</code> and then prints the contents as a string using the method <code>getContentAsString()</code>.

```
CREATE OR REPLACE MLE MODULE data_guide_mod
LANGUAGE JAVASCRIPT AS

import oracledb from "mle-js-oracledb";

export function createDataGuide() {

    // get a connection handle to the database session const connection = oracledb.defaultConnection();

    // get a SODA database const db = connection.getSodaDatabase();

    // open the collection const col = db.openCollection("employees_collection");
    if ( col === null ) {
        throw "'employees_collection' does not exist";
    }

    // generate a Data Guide const doc = col.getDataGuide();
    console.log(doc.getContentAsString());
}
```

7.3.16 Handling Transactions with SODA for In-Database JavaScript

Unlike the client-side JavaScript SQL driver, the MLE JavaScript SQL driver does not provide an autoCommit feature. You need to commit or roll your transactions back, either in the PL/SQL layer in case of module calls, or directly in the JavaScript code by calling connection.commit() or connection.rollback().



Caution:

If any uncommitted operation raises an error, and you do not explicitly roll back the transaction, the incomplete transaction might leave the relevant data in an inconsistent state (uncommitted, partial results).

7.3.17 Summary Example of Using SODA for In-Database JavaScript

An end-to-end example is included to demonstrate using SODA for In-Database JavaScript from start to finish.

Example 7-24 Use SODA for In-Database JavaScript

This example includes the following steps to use SODA for in-database JavaScript end-to-end:

- 1. Define a JavaScript module to interact with the employeesCol collection.
- 2. Create a JavaScript function to persist a JSON document in the collection.
- 3. Define another JavaScript function to list all the employees found in the collection.
- 4. Provide a call specification to each function within a PL/SQL package.

Define the JavaScript Module

This module makes use of the variables available in the global scope to limit the amount of coding required.

```
CREATE OR REPLACE MLE MODULE employees module
LANGUAGE JAVASCRIPT AS
/// <reference types="mle-js" />
const colName = "employeesCol";
* add an employee document to the employeesCol collection
^{\star} @param {json} employeeDoc - the new employee document to be added
export function addEmployee(employeeDoc) {
    // this code example uses the SODA database object (soda)
    // available in the global scope
    const col = soda.createCollection(colName);
    col.insertOne(employeeDoc);
}
/**
* return a JSON document containing all the employees from the collection
* @returns {json} a document listing employee id, last name, and
department name
*/
export function findAllEmployees() {
    // this array will contain all the employee documents
    // shortened to include employee ID, last name and
    // department name. It is safer to assume the array is
    // empty at first
    let employeeList = [];
    const col = soda.openCollection(colName);
    if ( col === null ) {
        return employeeList;
    }
    // iterate over all documents in the collection
```



```
const allDocs = col.find().getDocuments();
for ( let doc of allDocs ) {
    const payload = doc.getContent();
    // add the relevant fields to a JSON document before
    // pushing it to the employeeList array
    const employeeDetail = {
        "employee_id": payload.employee_id,
        "last_name": payload.last_name,
        "department_name": payload.department_name
    };
    employeeList.push(employeeDetail);
};

// return the list of employees found in the collection
    return employeeList;
}
```

Provide a Call Specification

Call specifications for both functions defined in the JavaScript module <code>employees_module</code> are provided by the following package. Note that the package is defined with invoker's rights rather than definer's rights. The invoker of the package must be granted the <code>SODA APP</code> role or else an error is thrown.

```
CREATE OR REPLACE PACKAGE employees_pkg
AUTHID current_user
AS

PROCEDURE add_employee(employee JSON)
AS MLE MODULE employees_module
SIGNATURE 'addEmployee';

FUNCTION find_all_employees
RETURN JSON
AS MLE MODULE employees_module
SIGNATURE 'findAllEmployees';

END employees_pkg;
/
```

Add Sample Data to the Collection

The following anonymous PL/SQL block converts two rows from $\mbox{HR.EMPLOYEES}$ to a JSON document and inserts them into the collection by invoking the JavaScript code externalized by the call specification created earlier.



```
'last name' IS e.last name,
            'email' IS e.email,
            'phone number' IS e.phone number,
            'hire date' IS e.hire date,
            'job id' IS e.job_id,
            'salary' IS e.salary,
            'commission pct' IS e.commission pct,
            'manager id' IS e.manager id,
            'department name' IS d.department name
    ТИТО
        1 employee
    FROM
        hr.employees e
        JOIN hr.departments d ON (e.department id = d.department id)
        e.employee id = 100;
    employees pkg.add employee(1 employee);
    SELECT
        JSON{
            'employee id' IS e.employee id,
            'first name' IS e.first name,
            'last name' IS e.last name,
            'email' IS e.email,
            'phone number' IS e.phone number,
            'hire date' IS e.hire date,
            'job id' IS e.job id,
            'salary' IS e.salary,
            'commission pct' IS e.commission pct,
            'manager_id' IS e.manager id,
            'department name' IS d.department name
    INTO
        1 employee
    FROM
        hr.employees e
        JOIN hr.departments d ON (e.department id = d.department id)
    WHERE
        e.employee id = 101;
    employees pkg.add employee(1 employee);
END;
```

Find All Employees in the Collection

Finally, EMPLOYEES_PKG.find_all_employees() can be used to provide a shortened list of employees stored in the collection.

```
SELECT

JSON_SERIALIZE(

employees_pkg.find_all_employees

PRETTY
```

```
) AS all_employees;
/
```

Result:

```
ALL_EMPLOYEES

"employee_id": 100,
    "last_name": "King",
    "department_name": "Executive"
},

{
    "employee_id": 101,
    "last_name": "Yang",
    "department_name": "Executive"
}

]
```



8

Post-Execution Debugging of MLE JavaScript Modules

The ability to easily debug code is central to a good developer experience. MLE provides the option to perform post-execution debugging on your JavaScript source code in addition to standard print debugging.

Post-execution debugging allows efficient collection of runtime state during program execution. Once execution of the code has completed, the collected data can be used to analyze program behavior and discover bugs that require attention. To perform post-execution debugging, you provide a debug specification that identifies the debugging information to be collected. A debug specification is a collection of debuggoints, each of which specify a location in the source code where debug information should be collected, as well as what information to collect. Debugpoints can be conditional or unconditional.

Note:

Post-execution debugging can only be applied to JavaScript code that is deployed as MLE modules. This debugging feature cannot currently be used when deploying code via dynamic execution.

Note:

MLE built-in modules such as the MLE JavaScript driver and MLE bindings cannot be debugged via post-execution debugging. An attempt to debug a built-in module will cause an ORA-04162 error to be raised.

For more information about MLE built-in modules, see Server-Side JavaScript API Documentation.

Module debugpoints apply to all executions of the module code, including via MLE call specifications, as well as via module import, whether from a dynamic MLE source or from another MLE module. Once enabled, a debug specification is active either until it is disabled or replaced by a new debug specification, or until the session ends.

Topics

- Specifying Debugpoints
 - Debugpoints are specified using a JSON document encoded in the database character set.
- Managing Debuggoints

Debugging can be enabled in a session by calling the procedure dbms mle.enable debugging with a debug specification.

Analyzing Debug Output

Output from debugpoints is stored in the Java Profiler Heap Dump version 1.0.2 format.

Error Handling in MLE

Errors encountered during the execution of MLE JavaScript code are reported as database errors.

8.1 Specifying Debugpoints

Debugpoints are specified using a JSON document encoded in the database character set.

Each debugpoint has the following elements:

- A *location* in the source code where the information is collected
- An action that describes what information to collect
- An optional *condition* that controls when debug information should be collected

Example 8-1 JSON Template for Specifying Debugpoints

```
{
at: <location-spec>,
action: [ <action-spec>, ... ],
[ condition: <condition-spec> ]
}
```

Debugpoint Locations

Debugpoint locations are specified via the line number in the source code of the application being debugged.

Debugpoint Actions

MLE post-execution debugging supports two kinds of actions: watch and snapshot.

Debugpoint Conditions

Both watch and snapshot can be controlled via conditions specified in the condition field.

8.1.1 Debugpoint Locations

Debugpoint locations are specified via the line number in the source code of the application being debugged.

The name of the MLE module to be debugged is specified via the name field and the location within the module where debug information is to be collected is specified via the line field. Example 8-4 provides an example JSON document with sample values.

8.1.2 Debugpoint Actions

MLE post-execution debugging supports two kinds of actions: watch and snapshot.

The watch action allows you to log the value of the variable named in the id field. The optional depth field provides you with control over the depth to which values of composite type variables are logged.

The snapshot action logs the stack trace at the point the snapshot action is invoked, along with the values of the local variables in each stack frame. A higher cost of performance is required by snapshot compared with watch but it provides a greater depth of information. As with the watch action, the optional depth field can be used to control the depth of logging for each variable. The depth parameter for the snapshot action applies to all variables captured by the action.

More precisely, the depth parameter controls how deeply you traverse the object tree in order to capture the value of a variable. For example, consider the following variable with nested objects:

```
let x = {
   a: {
      val: 42
},
b: 3.14
};
```

If the depth field is defined as 2, the object tree would be traversed and the value of the nested object a would be captured, which in this case is 42. If depth is specified as 1, the traversal would end at the first level, which would produce the following results:

```
x = {
"a": {
     "<unreachable>": true
};
"b": 3.14
}
```

The framesLimit field provides you with control over the number of stack frames to be logged. The default is to log all stack frames. framesLimit only applies to snapshot. Take, for example, a call hierarchy where a() calls b() and b() calls c(). If you take a snapshot in c(), framesLimit=1 would only capture the bottom-most stack frame (in this case, c()), framesLimit=2 would capture the bottom two (in this case, c()) and b()), and so on.

Example 8-2 JSON Template for Specifying Watch Action

To watch a variable, type must be set to watch. The id parameter is used to identify the variable or variables to watch and must be provided as either a string or an array of strings. The depth parameter is optional and is defined by a number.

```
actions: [
    { type: "watch",
        id: <string[]> | <string>,
        [depth : <number>] }
]
```



Example 8-3 JSON Template for Specifying Snapshot Action

To use the snapshot action, the type parameter must be set to snapshot. The framesLimit and depth fields are optionally provided as numbers.

```
actions: [
    { type: "snapshot",
        [framesLimit: <number>],
        [depth : <number>] }
]
```

8.1.3 Debugpoint Conditions

Both watch and snapshot can be controlled via conditions specified in the condition field.

The expression is evaluated in the context of the application at the location specified in the debugpoint and the associated action is triggered only if the expression evaluates to true.

There are no restrictions on the type of expression that can be included in the condition field. You must ensure that evaluating any expressions does not alter the behavior of the program being debugged.

Example 8-4 Watching a Variable in an MLE Module

The following code specifies a debugpoint for a module, myModule1, with two associated actions. A watch action for variable x with the logging depth restricted to 3, and a watch action for variable y with no restrictions on logging depth. The debugpoint also has an associated condition so that the debugpoint actions only trigger if the condition (x.id>100) is met.

```
{
    at : {
        name : "myModule1",
        line : 314
    },
    actions : [
        { type : "watch", id : "x", depth : 3 },
        { type: "watch", id : "y" }
    ],
    condition : 'x.id > 100'
}
```

8.2 Managing Debugpoints

Debugging can be enabled in a session by calling the procedure dbms mle.enable debugging with a debug specification.

In addition to an array of debugpoints, specified via the <code>debugpoints</code> field, a debug specification includes a version identifier, specified via the <code>version</code> field. The <code>version</code> field must be set to the value "1.0". Debug specifications can include debugpoints for multiple MLE modules.



Note:

Debug specifications require module names to be provided in the same case that they are stored in the dictionary. By default, module names are stored in uppercase unless the name is enclosed in double-quotation marks during module creation.

The procedure <code>dbms_mle.enable_debugging</code> also accepts a <code>BLOB sink</code> to which the debug output is written.

After the call to <code>dbms_mle.enable_debugging</code>, all debugpoints included in the debug specification are active. Every time one of the debugpoints is hit, the associated debug information is logged. The debug information is written out to the <code>BLOB sink</code> when control passes from MLE back to PL/SQL at the latest but could be written out in part or in full before this point:

- For dynamic MLE evaluations, control passes from MLE to PL/SQL when the call to dbms mle.eval returns.
- For MLE call specifications, control passes from MLE to PL/SQL when the call to the MLE call specification returns.

The installed debugpoints are active for all executions of the MLE modules regardless of which user's privileges the MLE code executes with.

Calling dbms_mle.enable_debugging again in the same session replaces the existing set of debugpoints. Debugpoints remain active until either the session ends or the user disables debugging explicitly by calling dbms mle.disable debugging.

Example 8-5 Enabling Debugging of an MLE Module

The debug specification in this example references the module <code>count_module</code>, created at the beginning of Example 5-4, and module in out example mod, created in Example 5-6.

```
DECLARE
  debugspec json;
  sink blob;
BEGIN
  debugspec:= json('
      "version": "1.0",
      "debugpoints": [
          "at": {
            "name": "COUNT MODULE",
            "line": 7
          },
          "actions": [
            { "type": "watch", id: "myCounter", depth: 1 }
          "condition": "myCounter > 0"
        },
          "at": {
            "name": "IN OUT EXAMPLE MOD",
            "line": 16
```



Debugging Security Considerations

Users must either own the MLE modules being debugged or have debugging privileges to it. This is necessary because the debugging feature allows you to observe runtime state of the MLE code.

COLLECT DEBUG INFO Privilege for MLE Modules

The COLLECT DEBUG INFO object privilege for MLE modules controls whether a user who does not own a module, but has EXECUTE privilege, can still perform debugging on said module.

8.2.1 Debugging Security Considerations

Users must either own the MLE modules being debugged or have debugging privileges to it. This is necessary because the debugging feature allows you to observe runtime state of the MLE code.

Additionally, because the <code>condition</code> field allows you to execute arbitrary code, this could potentially be used to alter the runtime behavior of the code being debugged. Concretely, you can use post-execution debugging on an MLE module if,

- You own the MLE module, or
- You have the COLLECT DEBUG INFO object privilege on the MLE module.

Privileges are checked every time code in an MLE module with one or more active debugpoints is executed. If you attempt to install debugpoints without the necessary privileges, an ORA-04164 error will be raised.

If an ORA-04164 is encountered, either

- The user who installed the debugpoints must be granted the COLLECT DEBUG INFO privilege on the module in question, or
- The debugpoints for the module must be disabled to continue executing code in the module in that session.



8.2.2 COLLECT DEBUG INFO Privilege for MLE Modules

The COLLECT DEBUG INFO object privilege for MLE modules controls whether a user who does not own a module, but has EXECUTE privilege, can still perform debugging on said module.

For instance, consider an MLE module, ModuleA, owned by user W. User W creates an invoker's rights call specification for a function in ModuleA and grants EXECUTE on this call specification on user V. For user V to have the ability to debug the code in ModuleA when calling this call specification, user W must also grant them the COLLECT DEBUG INFO privilege on ModuleA.

User $\ensuremath{\mathbb{W}}$ could use the following statement to grant user $\ensuremath{\mathbb{V}}$ the privilege to debug $\ensuremath{\mathtt{ModuleA}}$:

GRANT COLLECT DEBUG INFO ON ModuleA TO V;

The COLLECT DEBUG INFO privilege can subsequently be revoked if needed:

REVOKE COLLECT DEBUG INFO ON ModuleA FROM V;

8.3 Analyzing Debug Output

Output from debugpoints is stored in the Java Profiler Heap Dump version 1.0.2 format.

Every time a debugpoint is hit during execution, the debug information is saved as a heap dump segment. Once execution finishes, you have two options to analyze the debug output:

- Use the textual representation of the debug information obtained via the dbms mle.parse debug output function.
- Export the BLOB sink containing the debug output to an hprof file and use any of a number of existing developer tools to analyze the information.

Topics

- Textual Representation of Debug Output
 - The function <code>dbms_mle.parse_debug_output</code> takes as input a <code>BLOB</code> containing the debug information in the heap dump format and returns a JSON representation of the debug information.
- Analyzing Debug Output Using Developer Tools
 As an alternative to analyzing the textual representation of debug output, you also have the option to utilize tools such as JDeveloper, NetBeans, and Oracle Database Actions.

8.3.1 Textual Representation of Debug Output

The function <code>dbms_mle.parse_debug_output</code> takes as input a <code>BLOB</code> containing the debug information in the heap dump format and returns a JSON representation of the debug information.

The output of dbms_mle.parse_debug_output is an array of DebugPointData objects. DebugPointData represents the debug information logged every time a debugpoint is hit and comprises of an array of Frame objects. Each Frame includes the location in source code where the information was collected (the at field) and the names and values of local variables



logged at that location (the values field). Note that the keys of Frame.values are the names of the variables logged and the values are the values of those variables.

Example 8-6 demonstrates how you can specify a debugpoint in a sample JavaScript program and then use the function <code>dbms_mle.parse_debug_output</code> to produce a textual representation of the debug output.

Example 8-6 Obtain Textual Representation of Debug Output

The debugging shown later in this example is performed on the JavaScript function fib defined in the module fibunacci module:

```
CREATE OR REPLACE MLE MODULE fibunacci_module
LANGUAGE JAVASCRIPT AS
export function fib( n ) {
    if (n < 0) {
        throw Error("must provide a positive number to fib()");
    }
    if (n < 2) {
        return n;
    } else {
        return fib(n-1) + fib(n-2);
    }
}

CREATE OR REPLACE FUNCTION fib( p_value number)
RETURN NUMBER
AS MLE MODULE fibunacci_module
SIGNATURE 'fib(number)';
//</pre>
```

A debugpoint is placed at line 9 and then the <code>DBMS_MLE.PARSE_DEBUG_OUTPUT</code> function is used to view the debug information:

```
SET SERVEROUTPUT ON;
DECLARE
    1 debugspec JSON;
    l debugsink BLOB;
   l debuginfo JSON;
   1 value NUMBER;
BEGIN
    l debugspec := JSON ('
        version : "1.0",
        debugpoints : [
            {
                    name : "FIBUNACCI MODULE",
                    line : 9
                },
                actions : [
                    { type : "watch", id : "n" }
                ],
```



```
},
        ]
    }
    ');
    -- create a temporary lob to store the raw
    -- debug output
    DBMS LOB.CREATETEMPORARY( l debugsink, false );
    DBMS MLE.ENABLE DEBUGGING( 1 debugspec, 1 debugsink );
    -- run the application code
    l value := fib(4);
    DBMS MLE.DISABLE DEBUGGING;
    -- retrieve a textual representation of the debug
    -- output
    1 debuginfo := DBMS_MLE.PARSE_DEBUG_OUTPUT( 1 debugsink );
    DBMS OUTPUT.PUT LINE (
        json serialize(l debuginfo pretty)
    );
END;
Result:
[
      "at": {
        "name": "USER1.FIBUNACCI_MODULE",
        "line": 9
      "values": {
        "n": 4
  ],
      "at": {
        "name": "USER1.FIBUNACCI_MODULE",
        "line": 9
      "values": {
        "n": 3
  ],
  [
      "at": {
        "name": "USER1.FIBUNACCI_MODULE",
        "line": 9
```

8.3.2 Analyzing Debug Output Using Developer Tools

As an alternative to analyzing the textual representation of debug output, you also have the option to utilize tools such as JDeveloper, NetBeans, and Oracle Database Actions.

Once execution has finished, you can use the tool of your choice to inspect the values of local variables or to inspect the graph of variables at each point in time.

Integration with new tools can be developed as needed (e.g., Chrome Dev Tools) and UIs can be designed that are tailored specifically to the MLE use case.



Oracle Database Actions supports MLE post-execution debugging starting with Oracle Database release 23c, version 23.1.2.

See Also:

Using Oracle SQL Developer Web for more information about using Database Actions with MLE

8.4 Error Handling in MLE

Errors encountered during the execution of MLE JavaScript code are reported as database errors.

The database error raised depends on the type of error encountered. For example, syntax errors raise ORA-04160 while runtime errors (e.g., uncaught exceptions) raise

ORA-04161. The error message for each database error provides a brief description of the error encountered. Additionally, the DBMS_MLE PL/SQL package provides procedures to query the MLE JavaScript stack trace for the last error encountered in a dynamic MLE execution context or an MLE module in the current session.

The same security checks are made when calling <code>DBMS_MLE.get_ctx_error_stack()</code> as when calling <code>DBMS_MLE.eval()</code>. Thus, you cannot retrieve error stacks for MLE JavaScript code executing in dynamic MLE execution contexts created by other users.

DBMS_MLE provides a similar function, DBMS_MLE.get_error_stack(), to access the MLE JavaScript stack trace for application errors encountered during the execution of MLE modules. The function takes the module name and optionally the environment name as parameters, returning the stack trace for the most recent application error in a call specification based on the given arguments. If the module name or environment name is not a valid identifier, an ORA-04170 error is raised.

With MLE modules, it is only possible to retrieve the error stack for the module contexts associated with the calling user. This restriction avoids potentially leaking sensitive information between users via the error stack. A natural consequence of this restriction is that you cannot retrieve stack traces for errors encountered when executing definer's rights MLE call specifications owned by other users.

Example 8-7 Throwing ORA-04161 Error and Querying the Stack Trace

Executing the following code will throw an ORA-04161 error:

```
CREATE OR REPLACE MLE MODULE catch and print error stack
LANGUAGE JAVASCRIPT AS
export function f(){
    q();
}
function g() {
    h();
}
function h() {
    throw Error ("An error occurred in h()");
}
CREATE OR REPLACE PROCEDURE not getting entire error stack
AS MLE MODULE catch and print error stack
SIGNATURE 'f()';
BEGIN
    not getting entire error stack;
END;
```



Result:

```
BEGIN
ERROR at line 1:
ORA-04161: Error: An error occurred in h()
ORA-04171: at h (USER1.CATCHING AND PRINTING ERROR STACK:10:11)
ORA-06512: at "USER1.NOT GETTING THE ENTIRE ERROR STACK", line 1
ORA-06512: at line 2
You can query the stack trace for this error using the procedure
DBMS MLE.get error stack():
CREATE OR REPLACE PACKAGE get entire error stack pkg AS
    PROCEDURE get entire error stack;
END get entire error stack pkg;
CREATE OR REPLACE PACKAGE BODY get_entire_error_stack_pkg AS
    PROCEDURE print stack trace( p frames IN DBMS MLE.error frames t )
AS
    BEGIN
        FOR i in 1 .. p_frames.count LOOP
            DBMS OUTPUT.PUT LINE( p frames(i).func || '(' ||
            p frames(i).source || ':' || p frames(i).line || ')');
        END LOOP;
    END print stack trace;
    PROCEDURE do the work
    AS MLE MODULE catch and print error stack
    SIGNATURE 'f()';
    PROCEDURE get entire error stack AS
        l frames DBMS MLE.error frames t;
    BEGIN
       do the work;
    EXCEPTION
    WHEN OTHERS THEN
        l frames := DBMS MLE.get error stack(
            'CATCH AND PRINT ERROR STACK'
        print stack trace(l frames);
        raise;
    END;
END get entire error stack pkg;
BEGIN
    get entire error stack pkg.get entire error stack;
```



```
END;
```

The preceding code prints out the MLE JavaScript exception stack trace before raising the original error:

```
h (USER1.CATCH_AND_PRINT_ERROR_STACK:10)
g (USER1.CATCH_AND_PRINT_ERROR_STACK:6)
f (USER1.CATCH_AND_PRINT_ERROR_STACK:2)
BEGIN

*

ERROR at line 1:

ORA-04161: Error: An error occurred in h()

ORA-06512: at "USER1.GET_ENTIRE_ERROR_STACK_PKG", line 25

ORA-04171: at h (USER1.CATCH_AND_PRINT_ERROR_STACK:10:11)

ORA-06512: at "USER1.GET_ENTIRE_ERROR_STACK_PKG", line 11

ORA-06512: at "USER1.GET_ENTIRE_ERROR_STACK_PKG", line 11

ORA-06512: at "USER1.GET_ENTIRE_ERROR_STACK_PKG", line 18

ORA-06512: at line 2
```

Errors in Callouts

Database errors raised during callouts to SQL and PL/SQL via the MLE SQL driver are automatically converted to JavaScript exceptions.

Accessing stdout and stderr from JavaScript
 MLE provides functionality to access data written to standard output and error streams
 from JavaScript code.

8.4.1 Errors in Callouts

Database errors raised during callouts to SQL and PL/SQL via the MLE SQL driver are automatically converted to JavaScript exceptions.

For most database errors, JavaScript code can catch and handle these exceptions as usual. However, exceptions resulting from critical database errors cannot be caught. This includes:

- Internal database errors (ORA-0600)
- Fatal database errors (ORA-0603)
- Errors triggered due to resource limits being exceeded (ORA-04036)
- User interrupts (ORA-01013)
- System errors (ORA-7445)

Exceptions resulting from database errors that are either not caught or are re-signaled cause the original database error to be raised in addition to an MLE runtime error (ORA-04161). You can retrieve the JavaScript stack trace for such exceptions using DBMS MLE.get error stack() just like with other runtime errors.

8.4.2 Accessing stdout and stderr from JavaScript

MLE provides functionality to access data written to standard output and error streams from JavaScript code.

Within a database session, these streams can be controlled individually for each database user, MLE module, and dynamic MLE context. In each case, a stream can be:

- Disabled,
- Redirected to DBMS OUTPUT, or

module context.

- Redirected to a user provided CLOB
- Accessing stdout and stderr for MLE Modules
 The DBMS_MLE PL/SQL package provides the procedures set_stdout() and set stderr() to control the standard output and error streams for each MLE
- Accessing stdout and stderr for Dynamic MLE

 The procedures DBMS_MLE.set_ctx_stdout() and DBMS_MLE.set_ctx_stderr()

 are used to redirect stdout and stderr for dynamic MLE contexts.

8.4.2.1 Accessing stdout and stderr for MLE Modules

The <code>DBMS_MLE PL/SQL</code> package provides the procedures $set_stdout()$ and $set_stderr()$ to control the standard output and error streams for each MLE module context.

Alternatively, stdout can be redirected to DBMS_OUTPUT using the function DBMS_MLE.set_stdout_to_dbms_output(). The DBMS_MLE package provides an analogous function fore redirection stderr: DBMS MLE.set stderr to dbms output().

stdout and stderr can be disabled for a module at any time by calling DBMS MLE.disable stdout() and DBMS MLE.disable stderr() respectively.

By default, stdout and stderr are redirected to DBMS OUTPUT.

Note that the CURRENT_USER from an MLE function exported by the given MLE module may change depending on the CURRENT_USER when the function was called and whether the function is invoker's rights or definer's rights. A call to DBMS_MLE.set_stdout() or DBMS_MLE.set_stderr() by a database user, say user1, only redirects the appropriate stream when code in the MLE module executes with the privileges of user1.

In other words, one database user cannot ordinarily control the behavior of stdout and stderr for execution of an MLE module's code on behalf of another user.

All of these procedures take a module name and optionally an environment name as first and second arguments. This identifies the execution context whose output should be redirected. Omitting the environment name targets contexts using the base environment. Additionally, <code>set_stdout</code> and <code>set_stderr</code> take a user-provided CLOB as the last argument, specifying where the output should be written to.

Example 8-8 Redirect stdout to CLOB and DBMS OUTPUT for MLE Module

Consider the following JavaScript module:

```
CREATE OR REPLACE MLE MODULE hello_mod
LANGUAGE JAVASCRIPT AS
    export function hello() {
        console.log('Hello, World from MLE!');
    }
/
```



The following call specification makes the exported function hello() available for calling from PL/SQL code.

```
CREATE OR REPLACE PROCEDURE MLE_HELLO_PROC AS MLE MODULE hello_mod SIGNATURE 'hello'; /
```

The code below redirects stdout for the module hello_mod to a CLOB that can be examined later:

```
SET SERVEROUTPUT ON;
DECLARE
    l_output_buffer CLOB;
BEGIN
    -- create a temporary LOB to hold the output
    DBMS_LOB.CREATETEMPORARY(l_output_buffer, false);
    -- redirect stdout to a CLOB
    DBMS_MLE.SET_STDOUT('HELLO_MOD', l_output_buffer);
    -- run the code
    MLE_HELLO_PROC();
    -- retrieve the output buffer
    DBMS_OUTPUT.PUT_LINE(l_output_buffer);
END;
//
```

Executing the above produces the following output:

```
Hello, World from MLE!
```

Alternatively, stdout can be redirected to DBMS_OUTPUT using the function DBMS_MLE.SET_STDOUT_TO_DBMS_OUTPUT():

```
SET SERVEROUTPUT ON;
BEGIN
    DBMS_MLE.SET_STDOUT_TO_DBMS_OUTPUT('HELLO_MOD');
    MLE_HELLO_PROC();
END;
/
```

This produces the same output as before:

```
Hello, World from MLE!
```



8.4.2.2 Accessing stdout and stderr for Dynamic MLE

The procedures <code>DBMS_MLE.set_ctx_stdout()</code> and <code>DBMS_MLE.set_ctx_stderr()</code> are used to redirect <code>stdout</code> and <code>stderr</code> for dynamic MLE contexts.

The DBMS_MLE package similarly provides the procedures set_ctx_stdout_to_dbms_output() and set_ctx_stderr_to_dbms_output() to redirect stdout and stderr for dynamic MLE contexts to DBMS_OUTPUT.

A call to one of these functions redirects the appropriate stream for all dynamic MLE code executing within the context. However, any calls to MLE functions via the MLE SQL driver use the redirection effect for the MLE module that implement the function.

Example 8-9 Redirect stdout to CLOB and DBMS OUTPUT for Dynamic MLE

```
SET SERVEROUTPUT ON;
DECLARE
    l ctx DBMS MLE.context handle t;
    1 snippet CLOB;
    l output buffer CLOB;
BEGIN
    -- allocate the execution context and the output buffer
    l ctx := DBMS MLE.create context();
    DBMS LOB.CREATETEMPORARY(1 output buffer, false);
    -- redirect stdout to a CLOB
    DBMS MLE.SET CTX STDOUT(1 ctx, 1 output buffer);
    -- a bit of JavaScript code printing to the console
    1 snippet := 'console.log( "Hello, World from dynamic MLE!" )';
    -- execute the code snippet
    DBMS_MLE.eval(l_ctx, 'JAVASCRIPT', l_snippet);
    -- drop the execution context and print the output
    DBMS MLE.drop context(1 ctx);
    DBMS OUTPUT.PUT LINE(1 output buffer);
END;
```

This produces the following output:

Hello, World from dynamic MLE!

9

MLE Security

MLE utilizes a number of methods to support good security practices. This includes enforcing runtime state isolation, system and object privileges, and providing monitoring options.

Topics

- System and Object Privileges Required for Working with JavaScript in MLE
 Depending on the project's requirements, different privileges can be granted to users and
 or roles, allowing them to interact with JavaScript in the database.
- Security Considerations for MLE
 Besides the use of account privileges, MLE employs several other methods to ensure a
 high level of security.
- JavaScript Security Best Practices
 Details concerning the best practices when using features of MLE with JavaScript are described.
- MLE Security Examples
 Example scenarios are used to demonstrate security features used by MLE. The examples use a varying degree of separation between MLE modules, environments, and the necessary grants to enable the utilized functionality.

9.1 System and Object Privileges Required for Working with JavaScript in MLE

Depending on the project's requirements, different privileges can be granted to users and or roles, allowing them to interact with JavaScript in the database.

Administrators should review application requirements carefully and only grant the minimum number of privileges necessary to users. This is especially true for system privileges, which are very powerful and should only be granted to trusted users.

The minimum privilege required to work with MLE JavaScript code is the right to execute JavaScript code in the database. MLE distinguishes between dynamic MLE execution based on DBMS MLE and MLE execution using MLE modules and environments.

Creating stored code in JavaScript requires additional privileges to create JavaScript schema objects in your own schema.

The most powerful privileges available in MLE allow super-users to create, alter, and drop MLE schema objects in any schema, not just their own. As with all privileges in Oracle Database, those with ANY in their name are most powerful and should only be granted to trusted users if deemed absolutely necessary.

Note:

Object privileges on modules and environments do not grant access to an application, for example, the combination of source code and user context defined by a call specification (or through <code>DBMS_MLE</code>). This is achieved by granting access to the procedure or function object of the call specification.

See Also:

- Necessary Privileges for Creating MLE Modules and Environments in ANY Schema for more about handling system privileges
- Oracle Database Security Guide for more information about privileges in the Oracle Database

Topics

- Necessary Privileges for the Execution of JavaScript Code
- Necessary Privileges for Using the NoSQL API
- Necessary Privileges for Creating MLE Schema Objects
- Necessary Privileges for Creating MLE Modules and Environments in ANY Schema
- Necessary Privileges for Post-Execution Debugging

9.1.1 Necessary Privileges for the Execution of JavaScript Code

Before you can execute any JavaScript code in your own schema, the following object grant must have been issued to your user account:

```
GRANT EXECUTE ON JAVASCRIPT TO <role | user>
```

The EXECUTE ON JAVASCRIPT privilege does not include dynamic execution of JavaScript using <code>DBMS_MLE</code>. If you wish to make use of <code>DBMS_MLE</code>, an additional privilege is required:

GRANT EXECUTE DYNAMIC MLE TO <role | user>

9.1.2 Necessary Privileges for Using the NoSQL API

In cases where MLE JavaScript code references the Simple Oracle Document Access (SODA), the SODA APP role must be granted to the user or role:

```
GRANT SODA APP <role | user>
```



9.1.3 Necessary Privileges for Creating MLE Schema Objects

If you wish to create MLE modules and environments in your own schema, further system privileges are required:

```
GRANT CREATE MLE TO <role | user>
```

In case any MLE module is to be exposed to the database's SQL and PL/SQL layers in the form of call specifications, you also require the right to create PL/SQL procedures:

```
GRANT CREATE PROCEDURE TO <role | user>
```

It is highly likely that you will require further system privileges, depending on your use case, to create additional schema objects such as tables, indexes, and sequences. Beginning with Oracle Database 23c, the <code>DB_DEVELOPER_ROLE</code> role allows administrators to grant the necessary privileges to developers in their local development databases quickly, as shown in the following snippet:

```
BEGIN
        DBMS_DEVELOPER_ADMIN.GRANT_PRIVS('DEVUSER');
END;
/
```

See Also

Oracle Database Security Guide for more information about the DB_DEVELOPER_ROLE role

9.1.4 Necessary Privileges for Creating MLE Modules and Environments in ANY Schema

Additional privileges can be granted to power users and administrators, allowing them to create, alter, and drop MLE schema objects in *any* schema.

```
GRANT CREATE ANY MLE TO <role | user>
GRANT DROP ANY MLE TO <role | user>
GRANT ALTER ANY MLE TO <role | user>
```

As with all privileges in Oracle Databases featuring ANY in their name, these are very powerful and should only be granted after a thorough investigation to trusted users. For this reason, only the DBA role and the SYS account have been granted these privileges. The use of these system privileges is audited by the ORA SECURECONFIG audit policy.



To create MLE call specifications in schemas other than your own requires the right to CREATE ANY PROCEDURE to be granted as well:

GRANT CREATE ANY PROCEDURE TO <role | user>

Just like the previously listed system privileges, CREATE ANY PROCEDURE is audited by the same audit policy, ORA SECURECONFIG.



Oracle Database Security Guide for more information about the ORA SECURECONFIG audit policy

9.1.5 Necessary Privileges for Post-Execution Debugging

It is possible to allow other database users to collect debug information for MLE modules they don't own. By default, MLE owners can use post-execution debugging on their own MLE modules without specific grants. It is possible to grant the ability to collect debug information to a different role or user, allowing them to use post-execution debugging of JavaScript code on your behalf as the module owner:

GRANT COLLECT DEBUG INFO ON <module> TO <role | user>



You can elect to grant the execute privilege on MLE module calls created as PL/SQL code with definer's rights to users in other schemas. In this case, there is no need to grant other users any additional privileges.

Note:

Object privileges on modules and environments do not grant access to an application, for example, the combination of source code and user context defined by a call specification (or through <code>DBMS_MLE</code>). This is achieved by granting access to the procedure or function object of the call specification.

See Also:

Post-Execution Debugging of MLE JavaScript Modules for more information on post-execution debugging



9.2 Security Considerations for MLE

Besides the use of account privileges, MLE employs several other methods to ensure a high level of security.

Topics

- Initialization Parameter
- Execution Contexts
- Runtime State Isolation
- Database Security Model
- Considerations for Using MLE Call Specifications and Modules from Different Schemas
- Auditing MLE Operations in Oracle Database

9.2.1 Initialization Parameter

A new initialization parameter, MLE_PROG_LANGUAGES, allows administrators to enable and disable Multilingual Engine completely or selectively enable certain languages. It takes the values ALL, JAVASCRIPT, or OFF and it can be set at multiple levels:

- Container Database (CDB)
- Pluggable Database (PDB)
- Database session

If the parameter is set to OFF at CDB level, it cannot be enabled at PDB or session level. The same logic applies for PDB and session level: if MLE is disabled at the PDB level, it cannot be enabled at session level.



In Oracle Database 23c, MLE supports JavaScript as its sole language. Setting the parameter to ALL or JAVASCRIPT has the same effect.

Note:

Setting MLE_PROG_LANGUAGES to OFF prevents the execution of JavaScript code in the database, it does not prevent the creation or modification of existing code.

See Also:

Oracle Database Reference for more information about MLE PROG LANGUAGES



9.2.2 Execution Contexts

When executing JavaScript code in the database, MLE uses execution contexts to isolate runtime state such as global variables and other important information. Execution contexts are created implicitly when using modules and environments and explicitly when using DBMS MLE.

Regardless of the choice of JavaScript invocation, execution contexts are designed to prevent information leak.

The scope of JavaScript state never exceeds the lifetime of a database session. As soon as the session ends, either gracefully or forcefully, session state is discarded. If state needs to be preserved between sessions, you must persist it by storing it in a schema. If needed, state can be discarded by calling DBMS SESSION.reset package().



Oracle Database PL/SQL Packages and Types Reference for more information about DBMS SESSION

9.2.3 Runtime State Isolation

An MLE call specification is a PL/SQL unit referencing a function in an MLE module with an optional MLE environment attached. When you invoke a call specification in a session, the corresponding MLE module is loaded, the optional environment is applied, and the function specified in the call specification's signature clause is executed.

Before execution can begin, a corresponding execution context must be created (implicitly). Whether a new execution context is created or an existing context is reused depends on multiple factors, specifically:

- The MLE module referenced in the call specification
- The corresponding MLE environment
- The database user executing the call specification

Separate execution contexts are created to prevent information leak as well as undesired side effects such as global variables in a module being overwritten by accident.

With each invocation of a call specification, additional execution contexts are created. This is done so that modules cannot interfere with one another.

The main criteria for creating execution contexts in a user session are the MLE module name and the corresponding MLE environment. Call specifications referring to different combinations of MLE module and environment lead to different individual execution contexts being created.

Further separation between execution contexts is performed based on the user invoking the call specification.



Example 9-1 Runtime State Isolation Scenario

This example provides a sample scenario for runtime state isolation. Database user USER1 creates the following MLE schema objects:

```
CREATE OR REPLACE MLE MODULE isolationMod LANGUAGE JAVASCRIPT AS
let id;
             // global variable
export function doALotOfWork() {
  // a dummy function simulating a lot of work
  // the focus is on modifying a global variable
  id = 10;
}
export function getId() {
  return (id === undefined ? -1 : id)
}
CREATE OR REPLACE MLE ENV isolationEnv;
CREATE OR REPLACE PACKAGE context isolation package AS
  -- initialise runtime state
  procedure doALotOfWork as
    mle module isolationMod
    signature 'doALotOfWork()';
  -- access a global variable (part of session state)
  function getId return number as
    mle module isolationMod
    signature 'getId()';
  -- same function signature as before but referencing an environment
  function getIdwEnv return number as
   mle module isolationMod
    env isolationEnv
    signature 'getId()';
END;
/
When USER1, the owner of the MLE module, environment, and call specification (package),
calls context isolation package.doALotOfWork(), the global variable (id) is initialized to
10.
BEGIN
    context isolation package.doALotOfWork();
END;
```



Because <code>context_isolation_package.getId()</code> references the same MLE module and the same (default) environment as <code>context_isolation_package.doALotOfWork()</code>, the user's session has access to the global variable:

```
SELECT CONTEXT_ISOLATION_PACKAGE.getId;

GETID

10
```

When the combination of user, MLE module, and environment change, a new execution context is created. Although <code>context_isolation_package.getIdwEnv()</code> references the same MLE module as <code>getID()</code> and the user doesn't change, the function cannot retrieve the value of the global variable from the previously created execution context:

A value of -1 indicates that the global variable in the JavaScript module was found to be uninitialized.

If USER1, as the owner of the MLE call specification, grants the execute privilege on the package to another user, let's say USER2, a different execution context is created for USER2 even though the same function is called:

```
GRANT EXECUTE ON CONTEXT ISOLATION PACKAGE TO user2;
```

When USER2 tries to read the value of the ID, a new context is created and the return value indicating an uninitialized context is returned:

In this example, module and environment are identical between USER1 and USER2 as per the call specification. However, the fact that the function is called by a different user causes a new execution context to be created.

9.2.4 Database Security Model

The fewer privileges granted to program units, accounts, and roles, the less likely it is for them to be misused. As with every application, the principle of granting only the minimum number of necessary privileges should be followed. This is especially true in higher-tier environments like production. Technologies such as Privilege Analysis can be used to track down unnecessary privileges, allowing you to revoke them after careful regression testing.



Each MLE call specification is created within its own security context. The context includes information such as:

- The value of the AUTHID clause (definer or invoker)
- Whether or not privileges are inherited in invoker's rights calls
- Code Based Access Control
- Current user
- The qualified schema name
- Enabled Roles and Privileges in the absence of code based access control (CBAC) and invoker's rights

The combination of these attributes forms the security context of a code unit such as a MLE call specification or module. Note that no such security context exists for the JavaScript code stored in an MLE module.

PL/SQL allows you to easily change these attributes for each PL/SQL unit. A procedure can be executed with the invoker's rights or the definer's rights, roles can be attached to PL/SQL units, and cross-schema (execute) grants are commonplace. With each execution of a PL/SQL unit the security context may potentially change. This applies equally to MLE call specifications.

The situation is different with JavaScript code: the security context does not change for JavaScript-to-JavaScript calls. JavaScript functions do not have any notion of associated invoker's or definer's rights, or roles granted on the function itself. All of these apply only to (PL/SQL) call specifications.

JavaScript executed using <code>DBMS_MLE</code> is a little more strict when it comes to its security context. The combination of currently active user, roles/privileges, and schema in effect are recorded at the time the execution context is created by calling <code>DBMS_MLE.create_context()</code>. This combination must not change until the JavaScript code is executed and the context is removed, or else an error is thrown.



Oracle Database Security Guide for more information about Privilege Analysis

9.2.5 Considerations for Using MLE Call Specifications and Modules from Different Schemas

The same consideration that is used for other database applications written in, for example, PL/SQL apply for MLE JavaScript code as well. If a user is granted access to execute code from a schema other than their own, care needs to be taken to ensure the extent to which the code can use privileges of the calling user is appropriate.

Unlike PL/SQL, MLE JavaScript code stored in an MLE module is not associated with a particular set of roles, or any other notion of determining the security context in which the JavaScript code executes. From a high-level view, there are two important cases for cross-schema use of privileges:



- 1. USER1 invokes a call specification located in USER2's schema. The AUTHID clause of the call specification in USER2's schema determines whether the code owned by USER2's schema executes with the privileges of the invoker (USER1) or definer (USER2). In case of an invoker's rights call specification, potentially attached roles (CBAC) and the setting of INHERIT PRIVILEGES determine the active roles and privileges in addition to those granted by USER1 by roles or direct grants.
- 2. USER1 creates a call specification CallSpec_A for a module Module_A owned by USER1. CallSpec_A imports a JavaScript module Module_B owned by a different schema, USER2. The JavaScript code in Module_B is imported into an execution context created for USER1's call specification CallSpec_A. The JavaScript code in Module_B executes with the same privileges as any other JavaScript code in this execution such as in Module_A. USER1 must ensure that the code in Module_B is trustworthy and appropriate to execute with these privileges.

See Also:

Oracle Database Security Guide for more information about roles in definer's rights and invoker's rights PL/SQL units

9.2.6 Auditing MLE Operations in Oracle Database

Auditing is the monitoring and recording of configured database actions. As with any other auditable operations in Oracle Database, the use of MLE-related system privileges can be recorded.

Oracle provides the <code>ORA_SECURECONFIG</code> audit policy with the database. Starting with Oracle Database Release 23c, the audit policy includes the use of the following MLE system privileges:

- CREATE ANY MLE
- ALTER ANY MLE
- DROP ANY MLE

Administrators and security teams need to create and enable additional security policies if auditing the creation of MLE schema objects, including MLE modules, environments, and call specifications, is desired.

See Also:

Oracle Database Security Guide for more information about auditing in Oracle Database



9.3 JavaScript Security Best Practices

Details concerning the best practices when using features of MLE with JavaScript are described.

Topics

- Using Bind Variables for Security and Performance
- Generic Database and PL/SQL Specific Security Considerations
- Supply Chain Security
- · Software Bill of Material
- Using the Database to Store State
- Disabling Multilingual Runtime

9.3.1 Using Bind Variables for Security and Performance

The MLE JavaScript SQL driver allows you to use string concatenation to build SQL commands, including the predicates used in queries and DML statements. It is strongly recommended to avoid this bad practice as it is a major source for SQL injection attacks. Not only is the use of bind variables in SQL statements more secure than string concatenation but it is also more efficient as it allows the database to reuse the cursor in the shared pool.

If it is not possible to avoid the creation of dynamic SQL, ensure that you validate input to your code and scan for malicious content. The built-in <code>DBMS_ASSERT</code> package provides a wealth of functions designed to mitigate against SQL injection attacks. It does not offer complete protection but its use is very much recommended as it allows you to verify the following:

- The input string is a qualified SQL name
- The input string is an existing schema name
- The input string is a simple SQL name
- The input parameter string is a qualified SQL identifier of an existing SQL object

The use of bind variables for better security and scalability is not limited to a single programming language such as JavaScript, it equally applies to every development project using Oracle Database.

See Also:

- Server-Side JavaScript API Documentation for information about using bind variables with mle-js-oracledb
- Oracle Database Development Guide for more details regarding bind variables and their impact on performance and security



Example 9-2 Using Bind Variables Rather than String Concatenation

In this example, the SELECT statement accepts a bind variable rather than concatenation the input variable, managerID, to the SQL command.

```
CREATE OR REPLACE MLE MODULE select_bind LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function numEmployeesByManagerID(managerID) {
  const conn = oracledb.defaultConnection(managerID);
  const result = conn.execute(
    `SELECT count(*) FROM employees WHERE manager_id = :1`,
    [ managerID ]
  );
  return result.rows[0][0];
}
```

Example 9-3 Use DBMS_ASSERT to Verify Valid Input

In this example, the function <code>createTempTable()</code> creates a private temporary table to hold intermediate results from a batch process. The function takes a single argument: the name of the temporary table to be created (minus the prefix). The function checks if the parameter passed to it is a valid SQL name.

```
CREATE OR REPLACE MLE MODULE dbms assert module LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function createTempTable(tableName) {
  const conn = oracledb.defaultConnection();
 let result;
 let validTableName;
    result = conn.execute(
      `SELECT dbms assert.qualified sql name(:tableName)`,
      [tableName]
   );
   validTableName = result.rows[0][0];
  } catch (err) {
    throw (`'${tableName}' is not a valid table name`);
    return;
  result = conn.execute(
    `CREATE PRIVATE TEMPORARY TABLE ora\$ptt ${validTableName} (id
number) `
 );
}
```



If the table name passed to the function passes the test, it is then used to create a private temporary table using the default private_temp_table_prefix.

9.3.2 Generic Database and PL/SQL Specific Security Considerations

Because all JavaScript code is accessed eventually via a PL/SQL call specification, it is important to understand the implications of using PL/SQL as well. The following concepts are of particular importance:

- The difference between invoker's rights and definer's rights
- Code Based Access Control (CBAC)
- The impact of INHERIT PRIVILEGES in invoker's rights code
- Role grants and direct grants, both object as well as system privileges

You should always aim to only require the minimum security privileges (object and system) for JavaScript code to execute. This is especially important when you consider the use of external third-party JavaScript code.

Administrators should consider the use of encryption for both data at rest as well as data in motion.

See Also:

- Oracle Database Security Guide for more information about generic databaserelated security aspects
- Oracle Database Advanced Security Guide for information about encrypting data at rest using Transparent Data Encryption (TDE)

9.3.3 Supply Chain Security

Access to the rich community ecosystem is one of the advantages of using JavaScript in Oracle Database. Rather than creating functionality in-house and potentially duplicating effort, existing JavaScript can be used instead. While this is a convenient method for developing applications, it comes with certain risks.

In past years, the term supply chain attach has been used to describe the fact that certain popular open-source JavaScript modules have been abandoned by the original maintainers. Bad actors have taken some of these projects, becoming maintainers but only to inject malicious code into the source. The next time a project references such a compromised module, they incorporate the malicious code.

The same principles applied to client-side development apply to server-side development with MLE. Developers and security teams must be aware that code in the application executes with potentially elevated privileges. These can be abused by malicious code to compromise confidentiality, integrity, and availability properties of the application. For that reason, extra care must be taken to ensure third-party code is trustworthy and that the minimum number of privileges is granted to it. Many companies have a dedicated security team for vetting open-source modules prior to granting their approval to use them. At the very least, you should audit the JavaScript code that you are about to include in your project and document the result.



It is possible to lock a given version of an open-source module using a mechanism like the package-lock.json file so as not to get caught out if a new version of a module is distributed. Automatically pulling the latest version of an external code dependency is bad practice and should always be avoided.

In the case of JavaScript in MLE, JavaScript code executes with the database privileges that are in effect for the associated execution context. JavaScript code can retrieve and modify data stored in the database according to these privileges. Malicious code can leverage these privileges to modify the database in an inappropriate manner.

As a consequence, be sure to grant the privileges to create MLE modules carefully and only grant these in environments where they are essential. If possible, avoid granting the [CREATE | ALTER | MODIFY] ANY system privileges at all.

You should also review the INHERIT PRIVILEGES settings in the context of invoker's rights procedures. Once the settings for INHERIT PRIVILEGES are reviewed and secured according to industry best practice, consider the use of invoker's rights for MLE call specifications.

Additional higher levels of security for invoker's rights procedures can be achieved by implementing code based access control (CBAC). Using CBAC, developers can associate roles to PL/SQL units without having to elevate the privileges of the schema or invoker.



Oracle Database Security Guide for details about the INHERIT PRIVILEGES privilege

9.3.4 Software Bill of Material

Every project relying on external code in projects is strongly encouraged to maintain a record of all software components (including versions) that are bundled in a deployed application artifact.

The software bill of material (SBOM) is the key tool to use when reacting swiftly to a newly published vulnerability is of utmost importance. Exploits are almost guaranteed to be used immediately after a vulnerability has been published. Knowing exactly which version of a third-party library is in use allows you to save crucial time in preparing a response.

In addition to storing the actual code, MLE modules feature a metadata field that can be used to store arbitrary metadata with the module. In particular, it can be used to store an SBOM that describes all JavaScript libraries bundled in the module. The field is not interpreted by the MLE runtime. Content and format are entirely up to you.





MLE JavaScript Modules and Environments for more information about creating MLE modules and providing metadata to them

9.3.5 Using the Database to Store State

Applications written using MLE JavaScript code should not deviate from established patterns such as storing application state in tables. This allows you to make the best use of the rich number of security features available for Oracle Database.

In particular, you should not rely on JavaScript state that exceeds the boundaries of one stored procedure or function call.

Oracle Database has great support for JSON, offering both a relational as well as a NoSQL API. The database's JSON API is a natural candidate for MLE JavaScript code to store state. Storing state in Oracle Database provides a better programming model than application state, especially when it come to data persistence and transactional consistency.



Oracle Database JSON Developer's Guide for information about using JSON with Oracle Database

Example 9-4 Using Bind Variables Rather than String Concatenation

In this example, the SELECT statement accepts a bind variable rather than concatenation the input variable, managerID, to the SQL command.

```
CREATE OR REPLACE MLE MODULE select_bind LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function numEmployeesByManagerID(managerID) {
  const conn = oracledb.defaultConnection(managerID);
  const result = conn.execute(
    `SELECT COUNT(*) FROM employees WHERE manager_id = :1`,
    [ managerID ]
  );
  return result.rows[0][0];
}
```

Example 9-5 Use DBMS_ASSERT to Verify Valid Input

In this example, the function <code>createTempTable()</code> creates a private temporary table to hold intermediate results from a batch process. The function takes a single argument: the name of

the temporary table to be created (minus the prefix). The function checks if the parameter passed to it is a valid SQL name.

```
CREATE OR REPLACE MLE MODULE dbms assert module LANGUAGE JAVASCRIPT AS
import oracledb from "mle-js-oracledb";
export function createTempTable(tableName) {
 const conn = oracledb.defaultConnection();
 let result;
 let validTableName;
 try {
    result = conn.execute(
      `SELECT dbms assert.qualified sql name(:tableName)`,
      [tableName]
   );
   validTableName = result.rows[0][0];
  } catch (err) {
    throw (`'${tableName}' is not a valid table name`);
    return;
  result = conn.execute(
    `CREATE PRIVATE TEMPORARY TABLE ora\$ptt ${validTableName} (id
number) `
 );
}
```

If the table name passed to the function passes the test, it is then used to create a private temporary table using the default private temp table prefix.

9.3.6 Disabling Multilingual Runtime

In the case where a security vulnerability is detected in JavaScript code, you can prevent JavaScript code from execution by disabling the JavaScript runtime. Setting the initialization parameter <code>MLE_PROG_LANGUAGES</code> to <code>OFF</code> does not stop the database from accepting new code (such behavior prevents the implementation of a code fix) but it does stop anyone from executing JavaScript code.

Applications should be written with that option in mind. Once the MLE runtime is disabled, an error is thrown. Rather than showing the raw error to the end user, a more accessible error message should be created.

Although JavaScript does not have a specific lockdown feature, using the MLE_PROG_LANGUAGES parameter allows you to disable the MLE runtime at the session, PDB (lockdown profiles operate at this level), or CDB level. The COMMON_SCHEMA_ACCESS feature bundle in the lockdown profile can be used to disable MLE DDL.



9.4 MLE Security Examples

Example scenarios are used to demonstrate security features used by MLE. The examples use a varying degree of separation between MLE modules, environments, and the necessary grants to enable the utilized functionality.

Note that the examples are not fully usable on their own. The actual JavaScript code is not as important as the application's structure, such as:

- The schemas in which the code is located
- The call specification's syntax
- · The roles and privileges granted

Topics

- Business Logic Stored in MLE Modules
 - In this scenario, a user provides functionality implemented in JavaScript that is bound to a particular schema and relies on being executed as a particular user with certain privileges.
- Generic Data Processing Libraries
 In this scenario, generic JavaScript functionality is logically grouped inside a database schema. The JavaScript code is neither functionally nor logically tied to any existing
- Generic Libraries in Business Logic
 This scenario utilizes business logic contained in a single schema and extends functionality using generic libraries.

database objects. In other words, the processing logic is stateless.

9.4.1 Business Logic Stored in MLE Modules

In this scenario, a user provides functionality implemented in JavaScript that is bound to a particular schema and relies on being executed as a particular user with certain privileges.

This scenario covers the typical case of a back-end application centered around a single schema containing all necessary tables, indices, etc. Most importantly, the business logic is implemented as stored code in the database.

The JavaScript implementation in the form of MLE modules and an MLE environment is encapsulated in a single schema. Access to the functionality is only exposed using MLE call specifications based on one or multiple modules. Users of the application are granted execute privileges on (PL/SQL) call specifications only. No further privileges on MLE modules and environment are granted, nor are they necessary.

Consequently, the owner of the MLE modules controls access to the application through the AUTHID clause attached to the MLE call specifications. The pseudo-code in Example 9-6 demonstrates this scenario.

Example 9-6 Business Logic Stored in MLE Modules

In this example, the application schema is referred to as APP_OWNER . Note how MLE modules and environments are restricted to the APP_OWNER schema.

-- MLE Module containing helper functions commonly used by the application CREATE MLE MODULE app owner.helper module LANGUAGE JAVASCRIPT AS



```
export function setDebugLevel(level) {
 // ... JavaScript code ...
// ... additional functionality ...
-- An MLE Environment allowing other MLE Modules to import the helper
CREATE MLE ENV app owner.helper module env IMPORTS (
  'helperModule' module helper module
);
-- The main application module imports the helper module for common
CREATE MLE MODULE app owner.orders module LANGUAGE JAVASCRIPT AS
import { setDebugLevel } from "helperModule";
export function newOrder() {
setDebugLevel("INFO");
 // ... JavaScript code ...
export function delivery() {
  setDebugLevel("WARN");
  // ... JavaScript code ...
// ... additional functionality ...
-- The call specification is all the end users need to be granted
-- access to. The execute privilege to this definer's rights procedure
-- (created and executed with the app owner's database privileges)
-- is all that needs granting to the application role.
CREATE app owner.package orders pkg AS
  PROCEDURE new order AUTHID DEFINER AS
   MLE MODULE orders module
   ENV helper module env
    SIGNATURE 'newOrder()';
  PROCEDURE delivery AUTHID DEFINER AS
   MLE MODULE orders module
    ENV helper module env
    SIGNATURE 'delivery()';
END order pkg;
GRANT EXECUTE ON app owner.package orders pkg TO app role;
```

9.4.2 Generic Data Processing Libraries

In this scenario, generic JavaScript functionality is logically grouped inside a database schema. The JavaScript code is neither functionally nor logically tied to any existing database objects. In other words, the processing logic is stateless.

As there is no relation to any database schema objects such as tables or views, object grants are of no concern. The JavaScript code purely transforms functional arguments. Examples for such libraries include machine learning code, image manipulation like scaling, cropping, changes of resolution, etc. Other use cases include input validation or JSON processing.

The main purpose of the MLE modules deployed in such a fashion is to provide you with a common set of JavaScript tools that can be used in your own applications. Therefore, there aren't any pre-defined MLE call specifications provided. Instead, the schema containing these modules grants the execute privilege on MLE modules. It is up to the grantee to define MLE call specifications matching the use case. If necessary, MLE environments can be created alongside the MLE modules with respective grants to developers wishing to use the functionality created. Example 9-7 illustrates this scenario.

Example 9-7 Generic Data Processing Libraries

```
-- Common functionality potentially referenced by multiple applications
-- is grouped in a database schema. This particular MLE Module provides
-- input validation
CREATE MLE MODULE library owner.input validator module
  LANGUAGE JAVASCRIPT USING BFILE(js src dir, 'input validator.js');
-- Another MLE module provides common machine learning functionality
CREATE MLE MODULE library owner.commom ml module
  LANGUAGE JAVASCRIPT USING BFILE (js src dir, 'commom ml lib.js');
-- Rather than a Call Specification as demonstrated in Example 9-6,
-- this time the MLE Modules themselves are exported for use
-- in a different schema: frontend app
GRANT EXECUTE ON library owner.input validator module TO frontend app;
GRANT EXECUTE ON library owner.commom ml module TO frontend app;
-- frontend app makes explicit use of a select few functions exported
-- by the MLE modules
CREATE PACKAGE input validator pkg AS
  FUNCTION checkEMail (p email VARCHAR2) RETURN BOOLEAN AS
    MLE MODULE library owner.input validator module
    SIGNATURE 'checkEmail(string)';
  FUNCTION checkZIPCode (p zipcode VARCHAR2) RETURN BOOLEAN AS
    MLE MODULE library owner.input validator module
    SIGNATURE 'checkZIPCode(string)';
  -- additional functionality ...
END;
```



The grouping of common, stateless JavaScript code is not limited to a single schema. Further separation by feature, functionality, or maintainer is possible as well.

9.4.3 Generic Libraries in Business Logic

This scenario utilizes business logic contained in a single schema and extends functionality using generic libraries.

This example extends the scenarios demonstrated by Example 9-6 and Example 9-7. It is conceivable that the domain-specific business logic might require extension by common functionality such as logging or debugging. The latter can be written generically so that other applications can include it as well. There are numerous advantages to that approach including, but not limited to a unified framework for auxiliary functions.

In Example 9-8, the business logic in the APP_OWNER's schema, defined in Example 9-6, is extended with the previously introduced validation and machine learning functionality from Example 9-7.

There is no "best way" to work with MLE modules and environments in the database. It always depends on your particular use case. The included examples simply provide some background on how application logic can be grouped or separated, depending on a project's needs.

Example 9-8 Use Generic Libraries in Business Logic

```
-- Centrally managed JavaScript code library in the LIBRARY OWNER
schema
CREATE MLE MODULE library owner.commom ml module
  LANGUAGE JAVASCRIPT USING BFILE (js src dir, 'commom ml lib.js');
-- The grant makes the module available to APP OWNER
GRANT EXECUTE ON library owner.commom ml module TO app owner;
-- Business logic in schema APP OWNER makes use of the common ML
library
CREATE MLE MODULE app owner.helper module LANGUAGE JAVASCRIPT AS
export function setDebugLevel(level) {
  // ... JavaScript code ...
// ... additional functionality ...
-- A generic MLE environment references both APP OWNER's as well as
-- LIBRARY OWNER's MLE modules
CREATE MLE ENV app owner.all dependencies env imports (
  'helperModule' module helper module
  'commonML'
                module library owner.commom ml module
);
-- The main application module imports the helper module for common
-- as well as the common machine learning module provided by
```



```
LIBRARY OWNER
CREATE MLE MODULE app owner.orders module LANGUAGE JAVASCRIPT AS
import { setDebugLevel } from "helperModule";
import { churnRate }
                       from "commonML";
export function newOrder() {
  setDebugLevel("INFO");
  // ... JavaScript code ...
export function delivery() {
  setDebugLevel("WARN");
 // ... JavaScript code ...
export function estimateChurnRate() {
  // This function was imported from the common ML library
  // (an MLE module not stored in APP_OWNERs schema)
  const cr = churnRate();
  // ... JavaScript code ...
// ... additional functionality ...
-- the call specification is all the end-users need to be granted
-- access to. The execute privilege to this definer rights procedure
-- (created and executed with the app owner's database privileges)
-- is all that needs granting to the application role.
CREATE app owner.package orders pkg AS
  PROCEDURE new order AUTHID DEFINER AS
   MLE MODULE orders module
    ENV all dependencies env
    SIGNATURE 'newOrder()';
  PROCEDURE delivery AUTHID DEFINER AS
   MLE MODULE orders module
    ENV all dependencies env
    SIGNATURE 'delivery()';
  FUNCTION estimateChurnRate AUTHID DEFINER AS
   MLE MODULE orders module
    ENV all dependencies env
    SIGNATURE 'estimateChurnRate()';
END order pkg;
```

A

MLE Type Conversions

Supported conversions between JavaScript and PL/SQL, SQL, and JSON data types.

JavaScript target types include both native JavaScript types as well as SQL wrapper types. Supported SQL types are converted to the analogous JavaScript type by default where such a natural counterpart exists. If a conversion is attempted and there is no corresponding JavaScript type, conversion to a native JavaScript type is not supported and values are instead converted to the corresponding SQL wrapper type by default.

Note:

MLE does not provide functionality to prevent information loss that might occur between conversions from a customized database character representation to the built-in string representation of JavaScript (UTF-16).

See Also:

- Server-Side JavaScript API Documentation for information about using mle-js-bindings to change the default mappings when exchanging values between PL/SQL and JavaScript
- Server-Side JavaScript API Documentation for information on how to use mlejs-plsqltypes to create SQL wrapper types, such as OracleNumber
- Server-Side JavaScript API Documentation for information on using mle-jsoracledb to override the default conversions (as seen in Table A-1) when fetching column values from a SELECT statement

Date Conversions

JavaScript Date represents an instant (i.e., a single moment in time). Conversions can occur between the instant type Date and PL/SQL types DATE and TIMESTAMP that do not have time zone information. Conversions between instants on the JavaScript side and DATE and TIMESTAMP on the other side are handled as follows:

- When converting a Date to a TIMESTAMP or DATE, the instant is converted to a timezoneaware datatime value in the current session time zone. The local datatime portion of this value is stored in the target DATE or TIMESTAMP value.
- To convert a TIMESTAMP or DATE to a timezone-aware Date, the source datetime value is
 interpreted to be in the session time zone and is converted into an instant according to
 the session time zone.



Table A-1 Supported Mappings from SQL and PL/SQL Types to JavaScript Types

SQL Type	JavaScript Types (Bold Font Signifies Default)
NUMBER	number
	OracleNumber
BINARY_FLOAT	number
BINARY_DOUBLE	number
BINARY_INTEGER ¹	number
BOOLEAN	boolean
VARCHAR2	string
NVARCHAR2	string
CHAR	string
NCHAR	string
CLOB	OracleCLOB
	string
NCLOB	OracleCLOB
	string
BLOB	OracleBLOB
	Uint8Array (TypedArray)
RAW	Uint8Array (TypedArray)
DATE	Date
	OracleDate
TIMESTAMP	Date
	OracleTimestamp
TIMESTAMP WITH TIME ZONE	Date
	OracleTimestampTZ
TIMESTAMP WITH LOCAL TIME ZONE	Date
	OracleTimestampTZ
INTERVAL YEAR TO MONTH	OracleIntervalYearToMonth
INTERVAL DAY TO SECOND	OracleIntervalDayToSecond
NULL ²	null
JSON	object (JavaScript object or array) or corresponding JavaScript type for the JSON value ³

Note that BINARY_INTEGER is a PL/SQL type and not supported in SQL. MLE only supports BINARY_INTEGER on PL/SQL interfaces.



² Although not technically a type, MLE converts a SQL NULL value into a JavaScript null value and vice versa. This is so that JavaScript can indicate to the database that a value passed into the database is absent (for example, the return value of a function or an IN bind in a SQL statement).

³ See Table A-3 for details

Table A-2 Supported Mappings from JavaScript Types to SQL Types

JavaScript Type	SQL Type
number	NUMBER
boolean	
OracleNumber	
number	BINARY_FLOAT
number	BINARY_DOUBLE
number	BINARY_INTEGER
boolean	
number	BOOLEAN
OracleNumber	
boolean	
string	VARCHAR2
string	CHAR
string	NCHAR
string	NVARCHAR2
string	CLOB
OracleCLOB	
string	NCLOB
OracleCLOB	
string	UROWID
Uint8Array	BLOB
OracleBlob	
UintArray	RAW
Date	DATE
OracleDate	
Date	TIMESTAMP
OracleTimestamp	
Date	TIMESTAMP WITH (LOCAL) TIME ZONE
OracleTimestampTZ	
OracleIntervalYearToMonth	INTERVAL YEAR TO MONTH
OracleIntervalDayToSecond	INTERVAL DAY TO SECOND
null	NULL (any supported SQL type)



Table A-2 (Cont.) Supported Mappings from JavaScript Types to SQL Types

JavaScript Type	SQL Type	
number	JSON ²	
string		
boolean		
null		
undefined		
Date		
Uint8Array		
OracleNumber		
OracleDate		
OracleTimestamp		
OracleTimestampTZ		
OracleIntervalYearToMonth		
OracleIntervalDayToSecond		
object ¹		

- 1 JavaScript objects and arrays that do not match one of the classes listed above
- ² See Table A-4 for details

JSON Mapping Support

Values of the SQL $_{\rm JSON}$ type can be converted to and from JavaScript values. The type mapping between the SQL $_{\rm JSON}$ type and JavaScript values is aligned with type mappings employed by the $_{\rm node-oracledb}$ driver.



For more information about node-oracledb and the JSON data type, see the node-oracledb documentation.

Values of the SQL JSON type are converted to JavaScript values as follows:

- If the JSON value is an object, it is converted to an equivalent JavaScript object by converting all fields of the input object.
- If the JSON value is an array, it is converted to an equivalent JavaScript array by converting all elements of the input array.
- If the JSON value is a scalar value, it is converted to an equivalent value according to the type mapping in Table A-3.

Table A-3 Mapping from JSON Attribute Types and Values to JavaScript Types and Values

JSON Attribute Type or Value	JavaScript Type or Value
null	null
false	false



Table A-3 (Cont.) Mapping from JSON Attribute Types and Values to JavaScript Types and Values

JSON Attribute Type or Value	JavaScript Type or Value
true	true
NUMBER	Number
VARCHAR2	String
RAW	Uint8Array
CLOB	String
BLOB	UintArray
DATE	Date
TIMESTAMP	Date
INVERVAL YEAR TO MONTH	OracleIntervalYearToMonth
INTERVAL DAY TO SECOND	OracleIntervalDayToSecond
BINARY_DOUBLE	Number
BINARY_FLOAT	Number
Arrays	Array
Objects	A plain JavaScript Object

Values of a JavaScript type are converted to the SQL JSON type as follows:

- If the JavaScript value matches one of the scalar types in the first column of Table A-4, it is converted to a JSON value of the corresponding type.
- If the JavaScript value is an array, it is converted to a JSON array by converting all elements of the array. Note that <code>Uint8Array</code> values are treated as scalars as opposed to arrays, so <code>Uint8Array</code> values are converted to the type <code>RAW</code>, not to a JSON array.
- If the JavaScript value is an object that is neither an array nor matches any of the JavaScript types/ classes listed in Table A-4, it is converted to a JSON object. Each field of the object is converted according to the appropriate mappings.

Table A-4 Mapping from JavaScript Types and Values to JSON Attributes and Values

JavaScript Type or Value	JSON Attribute Type or Value
null	null
undefined	null
string	VARCHAR2
true	true
false	false
Uint8Array	RAW
Number	NUMBER
Date	DATE
OracleNumber	NUMBER
OracleDate	DATE
OracleTimestamp	TIMESTAMP
OracleTimestampTZ	TIMESTAMP WITH TIME ZONE



Table A-4 (Cont.) Mapping from JavaScript Types and Values to JSON Attributes and Values

JavaScript Type or Value	JSON Attribute Type or Value
OracleIntervalYearToMonth	INVERVAL YEAR TO MONTH
OracleIntervalDayToSecond	INTERVAL DAY TO SECOND
Array	Array
Object	Object



Index

A	See dropping a collection
	dictionary views
async/await interface, 5-10	USER_MLE_ENVS view, 2-20 USER_MLE_ENVS_IMPORTS view, 2-20
_	USER MLE MODULES view, 2-12
В	USER SOURCE view, 2-11
bind variables	Direct Fetch, 6-6
IN, 6-12	arrays, 6-7
IN OUT, 6-12	objects, 6-8
OUT, 6-12	documents
built-in modules, 6-25	creating, 7-12
accessing, 5-10	finding in collections, 7-20
mle-encode-base64, 2-10	inserting into collections, 7-15
mle-js-bindings, 2-10	removing from a collection, 7-27
mle-js-fetch, 2-10	replacing in collections, 7-26
mle-js-oracledb, 2-10	saving into collections, 7-17
mle-js-plsqltypes, 2-10	DRCP (database resident connection pool), 1-3
	dropping a collection, 7-11
С	dynamic execution, 1-5, 3-1
	workflow, 3-2
call specification, 1-6, 5-1, 5-11	
creating, 5-2	E
elements of, 5-4	
collections	ECMAScript, 1-8
creating, 7-7	available features, 6-1
dropping, 7-11	standard, 1-3, 1-4, 2-6
opening, 7-9	environment, 2-1
during creation, 7-7	execution context, 1-5, 1-8, 9-6
committing operations (transactions), 7-32	
creating a collection, 7-7	F
creating documents, 7-12	finding decomposite in callections 7.20
	finding documents in collections, 7-20
D	
Data Overal 0.4	G
Data Guard, 2-1	global variables F 10
data guide for a collection, getting, 7-31	global variables, 5-10
DBMS_MLE, 1-5	
EVAL procedure	Н
arguments of, 1-5 debugging	handling transactions 7.22
See post-execution debugging	handling transactions, 7-32
debugpoints, 1-9	
elements of, 8-2	I
specifying, 8-2	indexing decompose in a collection 7.00
deleting a collection	indexing documents in a collection, 7-29



initialization parameters	methods (continued)
MAX_STRING_SIZE, 1-3	save(), 7-17
MLE_PROG_LANGUAGES, 9-5	saveAndGet(), 7-17
inlined call specification, 5-1, 5-11	skip(), 7-19, 7-20
creating, 5-7	terminal and nonterminal, 7-19
elements of, 5-9	version(), 7-19, 7-20
inserting documents into collections, 7-15	MLE, 1-1-1-3
	call specification, 5-1
1	environment, 2-1
J	inlined call specification, 5-1
JavaScript, 1-3	module, 2-3, 4-1
global variables, 5-10	See also module
implementation of, 1-4	MLE JavaScript SQL driver, 2-10, 6-1
invoking, 1-5	selecting data, 6-6
loading from files, 3-3	MLE LANGUAGE clause, 5-7
providing inline, 3-2	MLE_PROG_LANGUAGES initialization
providing inline, 3-2	parameter, 9-5
	mle-js-fetch, 1-4
L	mle-js-oracledb
 	See MLE JavaScript SQL driver
language options, 2-17	module, <i>2-3</i>
js.console, 2-18	creating, 2-4
js.polyglot-builtin, 2-18	importing, 2-6, 4-1
js.strict, <i>2-18</i>	managing, 2-3
	naming, 2-3
M	module call, 1-6
	Multilingual Engine
MAX_STRING_SIZE initialization parameter, 1-3	See MLE
methods	OCC IVILL
count(), 7-19, 7-20	
create_context(), 9-8	N
createCollection(), 7-7	made made as a magnetic to the contract of the
createIndex(), 7-29	node package manager, 1-2
drop(), 7-11	nonterminal method
dropIndex(), 7-29	definition, 7-20
eval(), 1-5	NPM (node package manager), 1-2
execute(), 6-3	
	0
filter(), 7-19, 7-20	
find(), 7-20	opening a collection
getCollectionNames(), 7-10	during creation, 7-7
getContentAsString, 7-31	opening existing collections, 7-9
getCursor(), 7-19, 7-20	ORA SECURECONFIG audit policy, 9-10
getDataGuide, 7-31	= ' ' ' ' ' ' ' ' ' ' ' ' ' ' ' ' ' ' '
getOne(), 7-19, 7-20	D
hasNext(), 7-20	Р
headerOnly(), 7-19, 7-20	nest evecution debugging 1 0 9 2
key(), 7-19, 7-20	post-execution debugging, 1-9, 8-2
keys(), 7-19, 7-20	privileges, 1-4, 9-8
limit(), 7-19, 7-20	ALTER ANY MLE, 9-3
openCollection(), 7-9, 7-10	COLLECT DEBUG INFO, 9-4
read and write, 7-19	CREATE ANY MLE, 9-3
remove(), 7-19, 7-27	CREATE ANY PROCEDURE, 9-3
replaceOne(), 7-19, 7-26	CREATE MLE, 9-3
replaceOneAndGet(), 7-19, 7-26	CREATE PROCEDURE, 9-3
IESEL UALKAUEU. M-D	DROP ANY MLE, 9-3
reset_package(), 9-6	DROP ANY MLE, 9-3 EXECUTE DYNAMIC MLE, 9-2

privileges (continued) software bill of material (SBOM), 9-14 EXECUTE ON JAVASCRIPT, 9-2 INHERIT PRIVILEGES, 9-13 Т SODA_APP, 9-2 terminal method definition, 7-20 R transaction handling, 7-32 read methods, 7-19 removing documents from a collection, 7-27 U replacing documents in collections, 7-26 ResultSet object, 6-3, 6-6 Unicode Standard, 2-5 USER MLE ENV IMPORTS view, 2-20 S USER_MLE_ENVS view, 2-20 USER MLE MODULES view, 2-12 saving documents into collections, 7-17 USER_SOURCE view, 2-11 SBOM, 9-14 Simple Oracle Document Access (SODA), 7-1 W single-byte character set, 2-5 Smart-DB approach, 1-2 write methods, 7-19



SODA, 7-1