

Using Oracle® SQL Developer Web



Release 19.3

F22861-01

December 2019

The Oracle logo, consisting of the word "ORACLE" in white, uppercase, sans-serif font, centered within a solid red square.

ORACLE®

Using Oracle® SQL Developer Web, Release 19.3

F22861-01

Copyright © 2018, 2019, Oracle and/or its affiliates. All rights reserved.

Primary Authors: Celin Cherian, Mike Fitch

Contributing Authors: Chuck Murray

Contributors: Jeff Smith, Syme Kutz, Victor Martinez, Renato Uriel Gutierrez Salas, Antonio Hernandez, Sofia Isadora Padilla Munoz, Miguel Sanchez Villafan, Tulika Das, Eric Paapanen

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 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 installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. 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 and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon 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, Opteron, the AMD logo, and the AMD Opteron 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

Preface

Audience	v
Documentation Accessibility	v
Product Accessibility	v
Related Documents	v
Conventions	vi
Third-Party License Information	vi

1 About Oracle SQL Developer Web

1.1 About the SQL Developer Web User Interface	1-1
1.2 Accessing Oracle SQL Developer Web	1-3
1.3 Enabling User Access to SQL Developer Web	1-3

2 The Dashboard

3 The Home Page

4 The Activity Pages

4.1 The Activity Overview Page	4-1
4.2 The Automatic Indexing Page	4-1
4.3 The Performance Hub Page	4-3
4.4 The Alerts Page	4-5
4.5 The Sessions Page	4-5
4.6 The Storage Page	4-6

5 The Worksheet Page

5.1 Navigating Objects and Worksheets	5-1
5.2 Executing SQL Statements in the Worksheet Editor	5-2

5.2.1	Shortcut Keys in the Worksheet	5-4
5.3	Viewing the Worksheet Output	5-4
5.4	Uploading Data from Local Files	5-6

6 The Data Modeler Page

6.1	Navigating Diagrams and Objects	6-1
6.2	About the Data Modeling Editor	6-2

7 Creating and Editing Database Objects

7.1	The Table Properties Dialog	7-2
7.2	The Index Properties Dialog	7-12
7.3	The Sequence Properties Dialog	7-13
7.4	The View Properties Dialog	7-14
7.5	The Synonym Properties Dialog	7-15
7.6	Implied Foreign Keys	7-16

Preface

This online help provides information about Oracle SQL Developer Web, a browser-based interface for Oracle SQL Developer.

Audience

The online help is intended for those using the Oracle SQL Developer Web user interface. SQL Developer Web is available for users with access to Oracle Database services in Oracle Public Cloud.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Product Accessibility

Oracle SQL Developer Web provides features to support accessibility. See *SQL Developer Web Accessibility Guide*.

Related Documents

For the desktop version of Oracle SQL Developer, see *Oracle SQL Developer User's Guide*.

To download release notes, installation documentation, white papers, or other collateral for SQL Developer, go to the Oracle Technology Network (OTN) at

<http://www.oracle.com/technetwork/>

For the PL/SQL page on OTN, see <http://www.oracle.com/technetwork/database/features/plsql/>

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Third-Party License Information

Oracle SQL Developer contains third-party code. Oracle is required to provide the following notices. Note, however, that the Oracle program license that accompanied this product determines your right to use the Oracle program, including the third-party software, and the terms contained in the following notices do not change those rights.

colorpicker 1.1.2

MIT License Copyright (c) 2011 Martijn van der Lee Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Rappid 2.1

RAPPID GENERAL LICENSE STATEMENT AND LIMITED WARRANTY Copyright © 2009 - 2018 [client IO s.r.o.](#) Definitions "Rappid" shall mean a programmer's toolkit for building interactive diagramming applications as originally created by client IO. "Software" shall mean Rappid, as originally created by client IO. "Licensee" is you, only if you agree to be bound by the terms and conditions set forth in this Agreement. "Source code" means the JavaScript files and other related files in the Software package. This license statement and limited warranty constitutes a legal agreement ("License Agreement") between You (either as an individual or a single entity) and client IO for the software product ("Software") identified above, including any software, media, and accompanying on-line or printed documentation. By installing, copying, or otherwise using the Software, you are agreeing to be bound by the terms of this Agreement. If you do not agree to the terms of this Agreement, you are not authorized to use this Software. BY INSTALLING, COPYING, OR OTHERWISE USING THE

SOFTWARE, YOU AGREE TO BE BOUND BY ALL OF THE TERMS AND CONDITIONS OF THE LICENSE AGREEMENT. Upon your acceptance of the terms and conditions of the License Agreement, client IO grants to you the right to use the Software in the manner provided below. This license is perpetual, no additional payment is required to maintain it, with the exception of you breaking any part of this license, in which case you lose all rights under the license. You may transfer the rights granted to you under this agreement in and to the Software and documentation on a permanent basis provided you retain no copies, the recipient agrees to the terms of the License Agreement and the total number of developers licensed to use the software does not exceed the number of licensed developers at the time of license purchase. client IO must be informed in writing of the transfer. Except as provided in the License Agreement, you may not transfer, rent, lease, lend, copy, modify, translate, sublicense, time-share or electronically transmit or receive the Software, media or documentation. When transferring licenses, the whole of the licensed product must be transferred. In addition to the other terms contained herein, We grant to You a revocable, non-exclusive, non-transferable and non-sublicensable license to install and use the Software (the "Trial License") strictly for Your internal evaluation and review purposes and not for production purposes. This Trial License applies only if You have registered with Us for a Trial License of the Software and shall be effective for thirty (30) consecutive days following the date of registration ("the Trial Period"). You may only register for a Trial License once in any eighteen month period. You agree not to use a Trial License for any purpose other than determining whether to purchase a license to the Software. You are explicitly not permitted to distribute the Software to any user outside the Organization on whose behalf you have undertaken this license. Your rights to use the Trial License will immediately terminate upon the earlier of (i) the expiration of the Trial Period, or (ii) such time that You purchase a license to the Software. We reserve the right to terminate Your Trial License at any time in Our absolute and sole discretion. YOU ACKNOWLEDGE THAT TRIAL SOFTWARE MIGHT PLACE WATERMARKS ON OUTPUT, CONTAIN LIMITED FUNCTIONALITY, FUNCTION FOR A LIMITED PERIOD OF TIME, OR LIMIT THE FUNCTIONALITY OR TIME OF FUNCTIONING OF ANY OUTPUT. ACCESS TO AND/OR USE OF ANY FILES OR OUTPUT CREATED WITH SUCH SOFTWARE IS ENTIRELY AT YOUR OWN RISK. WE ARE LICENSING THE SOFTWARE ON AN "AS IS" BASIS AT YOUR OWN RISK AND WE DISCLAIM ANY WARRANTY OR LIABILITY TO YOU OF ANY KIND. DEVELOPER LICENSE TERMS To develop products using the Software, you require a number of development licenses equal to the number of engineers developing the application area that specifically uses the Software. This is the maximum number of such developers over any 1 month sliding window. A development license is not required in order to build an application that uses the software. A developer License may be installed on any number of computers at any time. client IO grants to you non exclusive license to install and use the Software for the sole purposes of designing, developing and testing application programs which you create. In addition to the right to install the complete software on the specified number of computers, the Source code may also be placed on storage devices for the purposes of version control, automated building and archiving. The entire downloaded Software package may also be archived for backup purposes. Additional Redistribution Terms for Software client IO hereby grants to you the worldwide, non exclusive, perpetual right under all intellectual property rights in and to the Software to incorporate the Software in whole or in part in your offerings and to redistribute the Software as incorporated to your customers and distributors and resellers. Rappid must be part of a larger system. You may not expose the Rappid API for developer use in your product. Updates Free updates are provided for all versions of the Software for the period of twelve (12) months from the time of the license purchase. Updates are performed by downloading the updated Software package through a link provided in the purchase process or sent to an email address provided

in the purchase process. Marketing client IO is permitted to reference you as a user of the Software in customer lists on any of the web sites owned by client IO, in presentations to clients and at trade events. Limited Warranty client IO warrants that as provided the Software does not contain any malicious code inserted to affect the functioning of the Software or any computer that the Software runs on. Any developer using Rappid must have a broad understanding of security issue relating to web application and must take all reasonable steps at a system level to mitigate those risks. client IO warrants that the Software does not infringe any patent, copyright or design secret of any third party. client IO is not aware of any facts upon which such a claim for infringement could be based. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLIENT IO AND ITS SUPPLIERS DISCLAIM ALL OTHER WARRANTIES AND CONDITIONS, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND NON-INFRINGEMENT, WITH REGARD TO THE SOFTWARE, AND THE PROVISION OF OR FAILURE TO PROVIDE SUPPORT SERVICES. THIS LIMITED WARRANTY GIVES YOU SPECIFIC LEGAL RIGHTS. YOU MAY HAVE OTHERS, WHICH VARY FROM STATE/JURISDICTION TO STATE/JURISDICTION. LIMITATION OF LIABILITY TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT SHALL CLIENT IO OR ITS SUPPLIERS BE LIABLE FOR ANY SPECIAL, INCIDENTAL, INDIRECT, OR CONSEQUENTIAL DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, OR ANY OTHER PECUNIARY LOSS) ARISING OUT OF THE USE OF OR INABILITY TO USE THE SOFTWARE PRODUCT OR THE PROVISION OF OR FAILURE TO PROVIDE SUPPORT SERVICES, EVEN IF CLIENT IO HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES AND JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF LIABILITY, THE ABOVE LIMITATION MAY NOT APPLY TO YOU. General Provisions This License Agreement may only be modified in writing signed by you and client IO. If any provision of this License Agreement is found void or unenforceable, the remainder will remain valid and enforceable according to its terms. If any remedy provided is determined to have failed for its essential purpose, all limitations of liability and exclusions of damages set forth in the Limited Warranty shall remain in effect. Governing Law and Jurisdiction This Agreement shall be subject to and governed by the Law of The Netherlands. Any dispute arising out of or in connection with this Agreement shall be exclusively dealt with by the courts of The Netherlands. This License Agreement gives you specific legal rights; you may have others which vary from state to state and from country to country. client IO reserves all rights not specifically granted in this License Agreement.

jQuery layout 1.4.3

This plug-in is dual-licensed under the GPL and MIT licenses.

GNU GENERAL PUBLIC LICENSE Version 3, 29 June 2007 Copyright © 2007 Free Software Foundation, Inc. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed. <http://www.gnu.org/licenses/gpl.html>

The MIT License Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in

all copies or substantial portions of the Software. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

CodeMirror 5.33

CodeMirror is an open-source project shared under an MIT license. MIT License Copyright (C) 2018 by Marijn Haverbeke <marijnh@gmail.com> and others Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Dexie.js 2.0.1

Dexie.js is licensed under the Apache License 2.0. Apache License Version 2.0, January 2004 <http://www.apache.org/licenses/> TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION 1. Definitions. "License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document. "Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License. "Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity. "You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License. "Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files. "Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types. "Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below). "Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from,

or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof. "Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution." "Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions: (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and (b) You must cause any modified files to carry prominent notices stating that You changed the files; and (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License. You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work\ by You to

the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions. 6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file. 7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License. 8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages. 9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability. END OF TERMS AND CONDITIONS

Hopscotch 0.3.1

Apache License Version 2.0, January 2004 <http://www.apache.org/licenses/> TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION 1. Definitions. "License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document. "Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License. "Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity. "You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License. "Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files. "Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types. "Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below). "Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial

revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof. "Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions: (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and (b) You must cause any modified files to carry prominent notices stating that You changed the files; and (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License. You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions

stated in this License. 5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work\ by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions. 6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file. 7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License. 8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages. 9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability. END OF TERMS AND CONDITIONS

1

About Oracle SQL Developer Web

Oracle SQL Developer Web is a browser-based application that uses ORDS (Oracle REST Data Services) to provide many of the database development and administration features of desktop-based Oracle SQL Developer.

Note:

Some features of Oracle SQL Developer Web are only available if you sign in as a user with database administration rights. For such features, a "restricted availability" statement appears at the start of the feature description; for example:



Available only if you signed in as a database user with administrator rights.

The main features include running SQL statements and scripts in the worksheet, exporting data, creating Data Modeler diagrams for existing schemas, enabling database administrators to monitor and manage the database, and providing database administrators with a real-time SQL monitoring interface

For a complete integrated development environment, see [Oracle SQL Developer](#).


1.1 About the SQL Developer Web User Interface

The SQL Developer Web user interface has three components:


- The [Header](#) at the top
- The page body, whose content varies depending on which page you are viewing
- The [Status Bar](#) at the bottom


Header

The header contains the selector icon, the main menu and the schema drop-down list.



- **Selector Icon**

Click the selector icon  to switch between the **Dashboard** and the **SQL Developer** headers.



- **Main Menu**

The main menu lists the pages available. The **Dashboard** is a single page and so has no main menu. The **SQL Developer** header's main menu has the following tabs:

- **Home:** Displays widgets that provide a general overview of the activity and status in SQL Developer Web. See [The Home Page](#).
- **Activity:** Provides several pages to view the performance and other characteristics of your database. See [The Activity Pages](#).
- **Worksheet:** Enables you to enter and execute SQL and PL/SQL commands. See [The Worksheet Page](#).
- **Data Modeler:** Enables you to create diagrams from existing database schemas, generate DDL statements, and create reports. See [The Data Modeler Page](#).


- **Schema Drop-Down List**

The schema drop-down list shows the database user you are signed in as, and provides the following items when you open it:

- **Preferences:** Opens a dialog that lets you set your preferred time zone.
- **Log:** Opens a dialog that shows the list of HTTP calls made during your SQL Developer Web session.
- **Help:** Opens online help for the page you are viewing in a new browser tab.
- **About:** Opens a dialog providing version information for the database and other components as well as copyright and licensing information.
- **Sign Out:** Signs you out of your database session.

Status Bar

The status bar contains icons that link to log files. The three icons (Errors, Warnings, Processes) are filters that have been applied to the log file.



Errors, Warnings: Displays an Errors or Warnings dialog, which lists log entries from unsuccessful REST calls or from any other problem in the application.

Processes: Displays a Processes dialog, which logs REST calls that are either finished or ongoing.

Log notification link: Displays a Log dialog, containing log entries of the following types: Errors, Warnings, Processes, SQL History and SQL Result.

1.2 Accessing Oracle SQL Developer Web

Oracle SQL Developer Web runs in Oracle REST Data Services and access to it is provided through schema-based authentication. To use Oracle SQL Developer Web, you must sign in as a database user whose schema has been enabled for SQL Developer Web.

In Oracle Autonomous Database databases, the ADMIN user is pre-enabled for SQL Developer Web. To enable another database user's schema, see [Enabling User Access to SQL Developer Web](#).

In the SQL Developer Web login page, enter the user name and the password of the database user for the enabled schema. Then, click **Sign In**. The Worksheet page is displayed.

1.3 Enabling User Access to SQL Developer Web

To enable a database user to sign into SQL Developer Web, run the following code as the ADMIN user:

```
BEGIN
  ords_admin.enable_schema(
    p_enabled => TRUE,
    p_schema => 'schema-name',
    p_url_mapping_type => 'BASE_PATH',
    p_url_mapping_pattern => 'schema-alias',
    p_auto_rest_auth => NULL
  );
  commit;
END;
```

where:

- *schema-name* is the database schema name in all-uppercase.
- *schema-alias* is an alias for the schema name that will appear in the URL the user will use to access SQL Developer Web. Oracle recommends that you do not use the schema name itself as a security measure to keep the schema name from being exposed.

After enabling user access, do the following:

1. Enter the following URL to go to the SQL Developer Web pre-login page:

```
http://<domain name>/ords/sql-developer
```

2. Enter the *schema-alias* and click **Next**. You see the URL change to `/ords/<schema alias>`.

3. In the SQL Developer Web login page, enter the password and click **Sign in**.

When you log out, you return to the pre-login page (`/ords/sql-developer`) and are prompted for the user name (*schema-alias*) again.

2

The Dashboard



Available only if you signed in as a database user with administrator rights.

The Dashboard displays general status information about the database and provides quick links to commonly used pages.

- **Database Status:** Displays the overall status of the database.
- **Alerts:** Displays the number of Error alerts in the alert log. You can click this item to open the Alerts page.
- **Database Storage:** Displays how much storage is being used by the database. You can click this item to open the Storage page.
- **Sessions:** Displays the status of open sessions in the database. You can click the title to open the Sessions page.
- **Physical IO Panel:** Displays the rates of physical reads and writes of database data.
- **Waits:** Displays how many wait events are occurring in the database by wait event class.
- **Quick Links:** Provides buttons to open the Worksheet, Data Modeler, and Performance Hub pages. It also provides a button to open the Oracle Application Express sign-in page for the current database.

3

The Home Page

The Home page contains widgets that provide a general overview of the activity and status in SQL Developer Web.

- **My Worksheets:** Displays the worksheets that have been created or modified. You can click the name of the worksheet to open it in the Worksheet page.
- **My Diagrams:** Displays the Data Modeler diagrams that have been created or modified. You can click the name of the diagram to open it in the Data Modeler page.
- **Recently Modified Objects:** Displays a timeline of the created, modified, and dropped objects in the database. You can zoom in and out using the + and – icons. You can also move horizontally by dragging the cursor to the right or left.
- **Invalid Objects:** Displays the invalid objects in the database.
- **Table Stats Freshness:** Displays the time period since the tables were last analyzed.

You can right-click the header in Invalid Objects, Table Stats Freshness, My Worksheets, or My Diagrams to manage or sort columns:

- **Columns:** Enables you to select columns to show or hide.
- **Sort:** Displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.

Right-click the body of the display table to count rows or to view records:

- **Count Rows:** Displays the number of rows in the table.
- **Single Record View:** Enables you to view data for a table or view, one record at a time.

4

The Activity Pages



Available only if you signed in to an Oracle Autonomous Database on dedicated Exadata infrastructure as a user with administrator rights.

The Activity menu on the SQL Developer Web header provides access to several pages to view the performance and other characteristics of your database.

- [The Activity Overview Page](#)
- [The Automatic Indexing Page](#)
- [The Performance Hub Page](#)
- [The Alerts Page](#)
- [The Sessions Page](#)
- [The Storage Page](#)

4.1 The Activity Overview Page



Available only if you signed in to an Oracle Autonomous Database on dedicated Exadata infrastructure as a user with administrator rights.

The Activity Overview page displays general information about the database.

- **Used Online Database Storage:** Displays how much storage is being used by the database. You can click the title to open the Storage page.
- **Users:** Displays how many user accounts are in the open, locked and expiring statuses. You can hover over one of the statuses to see a list of the user accounts with that status.
- **Sessions:** Displays the status of open sessions in the database. You can click the title to open the Sessions page.
- **Waits:** Displays how many wait events are occurring in the database for various reasons. You can click the title to open the Performance Hub page.

4.2 The Automatic Indexing Page



Available only if you signed in to an Oracle Autonomous Transaction Processing on dedicated Exadata infrastructure release 19c database or later as a user with administrator rights.

The automatic indexing feature automates the index management tasks in an Oracle database. Automatic indexing automatically creates, rebuilds, and drops indexes in a database thus improving database performance.

The Automatic Indexing page consists of the following parts:

- **Auto Index Settings:** Use this page to configure settings for automatic indexing in a database.
- **Activity Report:** Use this page to configure settings related to generating reports for automatic indexing operations in a database.

Auto Index Settings

In the Activity Report page, click the  (Open Settings) icon on the top right to configure settings for auto indexing. The Auto Indexing Settings page is displayed.

The configuration settings in the Auto Index Settings page are:

Indexing Mode

You can use the indexing mode configuration setting to enable or disable automatic indexing in a database. You can select one of the following values from the drop-down list:

- **IMPLEMENT:** New auto indexes are created as visible indexes and any existing invisible auto indexes are also set to visible indexes. In this mode, auto indexes are available to be used in SQL statements.
- **REPORT ONLY:** New auto indexes are created as invisible indexes and are not available to be used in SQL statements.
- **OFF:** The automatic indexing feature is disabled, that is, no new auto indexes are created and the existing auto indexes are disabled.

Retention Periods

- **Automatic Indexes:** Enter the number of days for which unused auto indexes are retained in the database, after which they are deleted. The default value is 373 days.
- **Manual Indexes:** Enter the number of days for which unused manually created indexes are retained in the database, after which they are deleted. When it is set to Default, the manually created indexes are not deleted by the automatic indexing process.
- **Reports:** Enter the number of days for which automatic indexing logs are retained in the database before they are deleted. As the automatic indexing report is generated based on these logs, the automatic indexing report cannot be generated for a period beyond this specified value. The default value is 31 days.

Schema Restrictions

You can specify schemas to include or exclude from using auto indexes. To specify the schemas to include, select **Include** and then in the Schema name text box, enter the schemas by selecting from the drop-down list. You can select multiple schemas. To remove a schema, click the x icon next to a schema name.

The automatic indexing process manages two schema lists, the inclusion list and the exclusion list. The inclusion list contains the schemas that can use auto indexes. The exclusion list contains the schemas that cannot use auto indexes. Initially, both these lists are empty and all the schemas in the database can use auto indexes when automatic indexing is enabled for a database.

To apply the configuration settings, click **Apply**. To return to the Activity Report page without applying the changes, click **Cancel**.

Activity Report

The report configuration settings in the Activity Report page are:

- **Report Interval Time:** Select the interval time of the report. You can select a value from the drop-down list or you can set a custom interval time.
- **Report Type:** Select the format of the report. The available options are HTML, TXT, JSON, XML.
- **Include Sections:** Select the sections to include in the report. By default, if you do not select any section to include, all sections will be included. The sections in the report are SUMMARY, INDEX_DETAILS, VERIFICATION_DETAILS, and ERRORS.
- **Report Level:** Select the level of automatic indexing information to include in the report. The values are BASIC, TYPICAL, and ALL.
- **Exclude Sections:** Select the sections to exclude from the report.

Click **Preview** to view a preview of the report based on the settings selected.

To download the report, click **Download** and the report is downloaded based on the format selected in the Report Type field.

4.3 The Performance Hub Page



Available only if you signed in to an Oracle Autonomous Database on dedicated Exadata infrastructure as a user with administrator rights.

The Performance Hub page shows performance data for a time period you specify. It consists of these parts:

- **Time Range Area.** Use the controls in time range area at the top of the page to specify the time period for which you want to view performance data.
- **ASH Analytics Tab.** Use this tab to explore ASH (Active Session History) information across a variety of different dimensions for the specified time period.
- **SQL Monitoring Tab.** Use this tab to view the top 100 SQL statement executions by different dimensions for the specified time period, and to view details of SQL statement executions you select.

Time Range Area

Use these controls in the time range area to specify the time period for which you want to view performance data:

- **Timeframe Dropdown.** Use this dropdown list, located in the top right of the time range area, to set the timeframe displayed in the timeline. You can choose **Last hour**, **Last 8 hours**, **Last 24 hours**, **Last Week**, or you can choose **Custom** and define your own timeframe.
- **Timeline.** The timeline displays a graph spanning the timeframe selected in the timeframe dropdown, showing Waits, User I/O and CPU usage during the period. At its end is the time slider.
- **Time Slider.** The time slider is a box you can drag back and forth along the current timeline. Use it to pick the specific time period within the timeframe for

which you want to view performance data. You can also drag the side handles on the time slider to make it wider or narrower to encompass a longer or shorter time period.

ASH Analytics Tab

The ASH Analytics tab consists of the Average Active Sessions chart and two secondary tables below it.

- **Average Active Sessions Chart.** This chart shows performance information for the time period defined by the time slider. You can choose to chart different dimensions of information, such as **Wait Class**, **Wait Event**, or **Service**, by selecting the dimension from the dropdown list to the right of the chart title.

You can download an AWR (Automatic Workload Repository) report for the current time period by right-clicking in the Average Active Sessions chart area and choosing **Generate AWR Report**.

For more information about ASH and AWR, see Active Session History (ASH) and Automatic Workload Repository (AWR) in *Oracle Database Concepts*.

- **Secondary Tables.** The two tables below the Average Active Sessions chart show the information dimension chosen in the chart filtered by another dimension you choose. For example, if the Average Active Sessions chart is showing Wait Class, you could show SQL ID and User Session dimension information by Wait Class, one in each of the two secondary tables.

In the secondary tables, the SQL ID and User Session dimensions provide links to SQL Details and Session Details pages, respectively, for the dimension data listed in the table.

SQL Monitoring Tab

The SQL Monitoring tab shows a table of the top 100 monitored SQL statements that were executing or that completed during the selected time period.

The table displays information about monitored SQL statement executions. If there is a green spinning icon in the Status column, then the monitored statement did not complete during the selected time period. A red cross indicates that the SQL did not complete either due to an error or due to the session getting terminated. If there is a check mark in the Status column, then the statement completed its execution during the selected time period.

SQL statements are monitored only if they have consumed at least 5 seconds of CPU or I/O time.

You can view information such as the status of a statement, its duration, its type (SQL, PL/SQL, or DBOP), its SQL ID, its SQL plan hash, the user who issued it, whether it executed as a serial or parallel statement, the time the database spent performing CPU activity, I/O, or other activity for the statement, the read and write requests and bytes associated with the statement, and the start and end time for the statement.

Click a SQL ID to display the SQL Details page with more information about that SQL statement.

4.4 The Alerts Page



Available only if you signed in to an Oracle Autonomous Database on dedicated Exadata infrastructure as a user with administrator rights.

The Alerts page is a chronological log of messages and errors and is commonly used to learn whether the background processes have encountered errors. You can review the alert log periodically to verify that your database system is operating normally.

The alert log includes the following:

- Nondefault initialization parameters used at startup
- Administrative operations, such as STARTUP, SHUTDOWN, ARCHIVE LOG, RECOVER, and CREATE/ALTER/ DROP DATABASE/TABLESPACE
- Messages and errors relating to the functions of certain background processes, such as LGWR
- Internal errors (ORA-600), block corruption errors (ORA-1578), and deadlock errors (ORA-60)

You can search for a specific value in the log by selecting the display column in the first drop-down list, selecting the condition in the second drop-down list, entering the search value in the box, and clicking the search icon.

In the display table, if you right-click the header row, you see:

Columns: Enables you to select columns to show or hide.

Sort: Displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.

If you right-click any other part of the display table, you see:

Count Rows: Displays the number of rows in the table.

Single Record View: Enables you to view data for a table or view, one record at a time.

4.5 The Sessions Page



Available only if you signed in to an Oracle Autonomous Database on dedicated Exadata infrastructure as a user with administrator rights.

The Sessions page shows information about all currently open sessions in the database. The data is automatically refreshed at intervals ranging from 10 seconds to 2 minutes. You can also refresh the data by clicking the Refresh icon at the top right of the screen.

The table shows summarized data about each open session. Select a session in the table to see more detailed data in the Session Details table below, such as the last SQL statement, explain plan, waits, contention, and so on. You can use the Column, Operator and Value fields to search for the required sessions.

In the display table, if you right-click the header row, you see:

Columns: Enables you to select columns to show or hide.

Sort: Displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.

If you right-click any other part of the display table, you see:

Count Rows: Displays the number of rows in the table.

Single Record View: Enables you to view data for a table or view, one record at a time.

4.6 The Storage Page



Available only if you signed in to an Oracle Autonomous Database on dedicated Exadata infrastructure as a user with administrator rights.

The Storage page shows the storage used based on the current allocation of tablespaces along with additional drill-down capabilities to view segments. You can refine the list of segments shown by using the filter feature. Click **View Datafiles** to view the datafiles in each tablespace.

You can view tablespace and segment space usage.

A tablespace is a database storage unit that groups related logical structures together. A tablespace is comprised of datafiles. A segment is a set of extents allocated from a tablespace for a specific database object such as a table or index.

To view space usage information

1. From the Database drop-down menu, click **Storage**.

The Storage page displays. If the Oracle database is version 12c or later, the Storage page shows the used and allocated storage space for tablespaces in any pluggable database. If the Oracle database is version 11g, the Storage page shows the used and allocated space for the entire database.

2. You can click a tablespace to view its storage information. An interactive report appears, showing the segments that exist within the tablespace. Most segments are user objects, and they include tables, LOBs, and indexes.
3. On the Segments page, you can refine the list of segments shown by using the filter feature.

For example, you can search for all the segments for a specific owner (schema) by selecting OWNER from the first drop-down list, entering the owner (schema) name in the box, and clicking the search icon.

5

The Worksheet Page

The Worksheet page provides a code editor that is used to enter and execute SQL and PL/SQL statements.

You can use SQL and PL/SQL statements in the worksheet to create a table, insert data, create and edit a trigger, select data from a table, and save that data to a file. Some other features are syntax highlighting and error detection.

Database administrators can also access the worksheet from the Quick Links pane in Dashboard.

The Worksheet page consists of the left pane for navigating worksheets and objects, the editor for executing SQL statements, and the output pane for viewing the results. These panes are described in the following sections:

- [Navigating Objects and Worksheets](#)
- [Executing SQL Statements in the Worksheet Editor](#)
 - [Shortcut Keys in the Worksheet](#)
- [Viewing the Worksheet Output](#)
- [Uploading Data from Local Files](#)



5.1 Navigating Objects and Worksheets


The Navigator tab in the left pane displays the saved objects for the selected schema. The drop-down menus enable you to select the schema for which you want to see the objects, and filter the results by object type.

The Worksheets tab displays the worksheets that are saved in the browser. The worksheets displayed are dependent on the browser used (Internet Explorer, Mozilla Firefox, or Google Chrome) for the SQL Developer Web session at the time of saving the worksheet, and they will not be accessible from any other browser. You can open or delete an existing worksheet by using the context (right-click) menu.

You can also search the contents of a saved worksheet or search for objects in the Navigator tab by name. The search functionality is not case-sensitive, retrieves all matching entries and does not require the use of wildcard characters.

Using the context (right-click) menu for objects in the Navigator tab, you can create and edit objects, add indexes and synonyms, and upload data from local files into existing tables.

To refresh the objects or worksheets list, click the  icon. To create a new object, click the  icon.

For context-related help, click the  icon.

You can drag objects from the Navigator and drop them into the worksheet editor in the right pane:

- If you drag and drop a table or view, you are prompted to select one of the following SQL statements: Insert, Update, Select, or Delete. For example, if you choose Select, a Select statement is constructed with all columns in the table or view. You can then edit the statement, for example, modifying the column list or adding a WHERE clause.

If you choose Object Name, the name of the object prefixed by the schema name is added to the worksheet.


- If you drag and drop a function or procedure, you can choose to insert the name or the PL/SQL code of the function or procedure in the worksheet. If you select the PL/SQL code, you can enter the parameters before inserting the code into the worksheet.

5.2 Executing SQL Statements in the Worksheet Editor


The worksheet editor in the right pane enables you to enter the SQL statements that you intend to execute. You can use SQL and PL/SQL statements to specify actions such as creating a table, inserting data, selecting data or deleting data from a table. For multiple statements, each non-PL/SQL statement must be terminated with either a semicolon or (on a new line) a slash (/), and each PL/SQL statement must be terminated with a slash (/) on a new line. SQL keywords are automatically highlighted.

For a list of the SQL*Plus statements that are supported by the SQL Worksheet, see Supported SQL*Plus Statements.

If you press **Ctrl+Space**, the worksheet provides you with a list of possible completions at the insertion point that you can use to autocomplete code that you are editing. This list is based on the code context at the insertion point. Also, you can select multiple options in the list using **Ctrl+Click**.



An error in the code is signified by a red dot or squiggle line. When you hover over it, you see a pop-up displaying possible fixes for resolving the error.



The SQL Worksheet toolbar contains icons for the following operations:

- **Worksheet** enables you to create a worksheet, open or delete an existing worksheet, and save a worksheet.

 **Note:**

The worksheets are saved in the browser. Therefore, if you do sensitive work on a computer in a public network, remember to clear the browser cache before you leave. Additionally, the saved worksheets will only be displayed in the browser that was used at the time of creating the worksheet.

- **Run Statement** executes the selected statements or the statement at the mouse pointer in the worksheet editor. The SQL statements can include bind variables and substitution variables of type VARCHAR2 (although in most cases, VARCHAR2 is automatically converted internally to NUMBER if necessary). A dialog box is displayed for entering variable values.
- **Run Script** executes all statements in the worksheet editor using the Script Runner. The SQL statements can include bind variables (but not substitution variables) of type VARCHAR2 (although in most cases, VARCHAR2 is automatically converted internally to NUMBER if necessary). A dialog box is displayed for entering bind variable values.
- **Explain Plan** generates the execution plan for the statement (internally executing the EXPLAIN PLAN statement). The execution plan is automatically displayed in the Explain Plan tab in the worksheet output pane.
- **Autotrace** runs the statement and collects runtime statistics and the actual execution plan. The Autotrace output is displayed in the Autotrace tab in the worksheet output pane. Currently, there are no preferences available.
- **Download Editor Content** downloads the content of the worksheet as a SQL file to the local system.
- **Format** formats the SQL statement in the editor, such as capitalizing the names of statements, clauses, keywords, and adding line breaks and indentation.
- **Clear** removes the statements from the editor.

- **Tour** provides a guided tour of the worksheet highlighting salient features and providing information that is useful if you are new to the interface.
- **Help** provides context-related help and provides a link to the help documentation.

For a list of shortcuts used in the worksheet, see [Shortcut Keys in the Worksheet](#).

Related Topics

- Execution Plan in *Oracle SQL Developer User's Guide*
- Autotrace Pane in *Oracle SQL Developer User's Guide*

5.2.1 Shortcut Keys in the Worksheet

The following table lists the shortcut keys for various user actions in the worksheet.

Table 5-1 Shortcut Keys for User Actions in the Worksheet

Shortcut	Action
Ctrl+Enter/ Cmd+Enter	Runs the code as query.
Ctrl+Down Arrow/ Cmd +Down Arrow	Moves to the next SQL code from history.
Ctrl+Up Arrow/ Cmd +Up Arrow	Moves to the previous SQL code from history.
Ctrl+D/ Cmd +D	Clears the editor.
Ctrl+S/ Cmd +S	Saves the current worksheet.
Ctrl+O/ Cmd +O	Opens the worksheet browser dialog.
Ctrl+I/ Cmd+I	Downloads the content of the editor.
F1	Opens the help topic.
F5	Runs code as script.
F6	Shows Autotrace.
F10	Shows Explain Plan.
Ctrl+F7/ Cmd +F7	Formats code in the editor.
Ctrl+Space/ Cmd+Space	Autocompletes code (shows hints).
Windows +Esc/ Cmd +Esc	Focuses outside the editor and navigates to the rest of the application using the Tab key.

5.3 Viewing the Worksheet Output

The bottom right pane of the Worksheet screen has tabs that display the following panes:

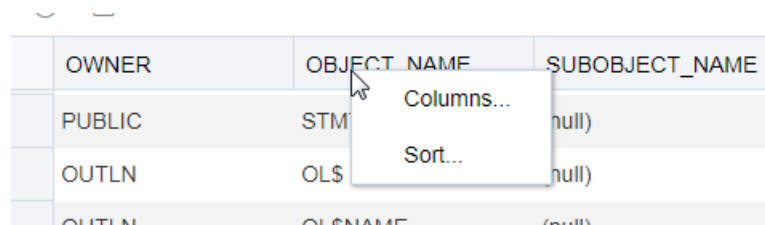
- **Query Result:** Displays the results of the most recent Run Statement operation in a display table.
- **Script Output:** Displays the text output from your statements executed as a script using SQL Developer Web's script engine.
- **DBMS Output:** Displays the output of DBMS_OUTPUT package statements.
- **Explain Plan:** Displays the plan for your query using the Explain Plan command.
- **Autotrace:** Displays the session statistics and execution plan from `v$sql_plan` when executing a SQL statement using the Autotrace feature. Displays the output if you clicked the Autotrace icon.
- **SQL History:** Displays the SQL statements and scripts that you have executed. To re-enter a previously executed query in the worksheet, double-click the query in the history list. You can search for specific statements by clicking the Search icon. The Search functionality is case-sensitive, retrieves all entries that contain the search text, and does not require wildcard characters.

The icons in this pane are:

- **Clear output:** Clears the output.
- **Show info:** Displays the SQL statement for which the output is displayed.
- **Open in new tab:** Opens the query result or explain plan in a new window.
- **Download:** This is applicable only for Query Result. Enables you to download the query result to your local computer in CSV, JSON, XML, or TEXT (.tsv) format.

In the Query Result tab, in the display table, the context menu (right-click) for the row header consists of the following:

- **Columns** enables you to select columns to hide.
- **Sort** displays a dialog box for selecting columns to sort by. For each column, you can specify ascending or descending order, and you can specify that null values be displayed first.



OWNER	OBJECT_NAME	SUBOBJECT_NAME
PUBLIC	STM	(null)
OUTLN	OL\$	(null)
OUTLN	OL\$NAME	(null)

The context menu for the rest of the display table consists of the following commands:

- **Count Rows** displays the number of rows in the result set for your query.
- **Single Record View** enables you to view data for a table or view, one record at a time.
- **Export** generates the file for download based on the format selected, which can be XML, CSV (comma-separated values including a header row for column identifiers), Insert , Delimited, Fixed, HTML, JSON, or TEXT.

 **Note:**

If a popup blocker is enabled, it will prevent the file from downloading.

	OBJECT_NAME	SUBOBJECT_NAME	OB...
	STMT_AUDIT_OPTIO	(null)	
	OL\$		
	OL\$NAM		
	OL\$SIG		
	OL\$HIN		
	OL\$HNT NIM	(null)	

Related Topics

- DBMS Output Pane in *Oracle SQL Developer User's Guide*
- Execution Plan in *Oracle SQL Developer User's Guide*
- Autotrace Pane in *Oracle SQL Developer User's Guide*

5.4 Uploading Data from Local Files

In the Worksheet page, you can upload data from local files into an existing table. The file formats that you can upload are CSV, XLS, XLSX, TSV and TXT.


To upload data from a local file to an existing table:

1. In the Navigator tab, in the left pane, right-click the table that you want to upload data into, select **Data Loading**, and then select **Upload Data**.

The Import data dialog is displayed.

2. Drag and drop the file from your system into the dialog or click **Select Files** to browse for the file and open it.

A preview of the data is displayed in a grid format.

3. Click the  (Show/Hide options) icon to display options that you can modify for data preview:

- **Column names:** Select **First row** to display column headers in the first row.
- **Text enclosure:** This option is visible only when the selected file is in a plain text format (CSV, TSV, or TXT). Select the appropriate separator to part the data in the file.
- **Rows to skip:** Use the up and down arrows to select the number of rows to skip.
- **Rows to load:** Use the up and down arrows to select the number of rows to load.
- **Preview size:** Use the up and down arrows to select the number of rows to preview.

To remove the options selected and the data preview, click **Clear**.

After selecting the required options, click **Apply**, and then click **Next**.

4. In Column mapping, match the data in the file to the appropriate columns in the target table. By default, SQL Developer Web will match using the column name.


To modify, click the Show/Hide options icon. In Match columns by:

- Select **Name** to match columns based on the name of the column in the target table.
- Select **Position** if you want to match columns based on the position of the column in the target table.
- Select **None** to remove the current selections and to select the target column for each source column from the drop-down list.

 **Note:**

SQL Developer Web will attempt to retrieve the correct format mask of date-based columns based on the data in the file, but it will still enable you to change the suggested format by entering the required format directly in the target column.

If there are any issues to be resolved, you see a notification such as

 **1 pending actions** on the top right that provides details about the issue to be fixed.

Click **Next**.

5. A summary of the previous screens is displayed. Click **Finish**.

The data will start uploading to the target table. After it is completed, an entry is added to the Log with the status of the operation. To view the Log, click the timestamp notification at the bottom of the page. If the operation is successful, a `Data Import Completed` notification is displayed.

6. For a detailed summary of the upload process, right-click the table in the Navigator tab, select **Data Loading**, and then select **Loaded Data**. A summary of the data loaded is displayed in the Loaded data summary dialog.

If any data failed to load, you can view the number of rows in the Failed Rows column. Click the column and a dialog is displayed showing the failed rows.

In the Loaded data summary dialog, you can also search for files loaded by schema name, table name, or file name. To remove the loaded files, click the Delete icon.

6

The Data Modeler Page

The Data Modeler page provides an integrated version of Oracle SQL Developer Data Modeler with basic reporting features. You can create diagrams from existing schemas, retrieve data dictionary information, generate DDL statements, and export diagrams.

 **Note:**

You can create and edit objects for your schema in the Data Modeler interface.


- [Navigating Diagrams and Objects](#)
- [About the Data Modeling Editor](#)

Related Topics

- *Oracle SQL Developer Data Modeler User's Guide*

6.1 Navigating Diagrams and Objects

The **Diagrams** tab lists the Data Modeler diagrams that have been saved. When you right-click a diagram, you have options to open, save, delete, and view properties. When you click a diagram in the list, a thumbnail preview is shown with an option to open the diagram.

- To create a new diagram, in the Diagrams tab, click the  icon next to the Search field.
- To open an existing diagram, in the Diagrams tab, select or right-click the diagram, and then click **Open**.

 **Note:**

Diagrams are stored in the database, in the SQL Developer Web user schema. When you use Data Modeler for the first time, a `OSDDMW_DIAGRAMS` table is created in your schema, to store the diagrams.

Oracle SQL Developer Data Modeler can import from or export to the `OSDDMW_DIAGRAMS` table if a connection is provided. For more information, see [Sharing Diagrams with SQL Developer Data Modeler](#).

The **Navigator** tab lists the objects that are available for each schema. You can select the schema and object type from the drop-down lists. For a particular schema, if a

tables object is selected, the left pane displays all the tables that belong to the particular schema. You can expand a table to view its columns. You can drag and drop an object into a selected diagram in the right pane. You can also add objects to the diagram by clicking the Add Objects to Diagram icon in the editor pane.

You can also search the contents of a saved diagram or search for objects in the Navigator tab by name. The search functionality is not case-sensitive, retrieves all matching entries and does not require the use of wildcard characters.

The context menu for a table or view consists of:

Add Object to Diagram: Adds the selected object to the selected diagram.

Add Object with dependencies to Diagram: Adds parent and child tables related to the object to the selected diagram.

Add Object as Star Schema to Diagram: Adds the selected object to the diagram and searches the data dictionary for foreign keys and implied foreign keys related to the object. The related tables or views are added to the diagram and the star schema layout is applied. See [Implied Foreign Keys](#)

New, Edit, Use as Template: Creates a new object (such as table or view) for a specific schema, edits an existing object, or creates a new object by using an existing one for the initial content. See [The Table Properties Dialog](#) and [The View Properties Dialog](#).

Add/Edit Sequence: To create or edit a sequence for a selected schema. See [The Sequence Properties Dialog](#).

Add Index: To create an index for a table. See [The Index Properties Dialog](#).

6.2 About the Data Modeling Editor

You can create and work with relational diagrams in the editor pane.

When an object is selected, you can inspect the properties of the object in the right pane. The only properties that you can edit are background and border color. For example, when a table is selected, the properties displayed are Columns, Primary Key, Unique Constraints, Indexes, Foreign Keys, Comments, and Colors. You can edit the background and border color by selecting from the color palette in the Pick a Color dialog box. Click the Reset icon next to the Background and Border Color fields to return to the default colors.

When creating diagrams, you can do the following:

- Move objects around by dragging the mouse.
- Adjust or move objects with the relationships intact.
- Add elbows to relationship lines to avoid intersecting with lines from other objects. Right-click the relationship line and drag to create the elbow. Click the X icon to restore to the original shape.
- Resize objects by dragging the handles that are positioned around the box.
- Right-click an object for options to preview DDL statements, update, delete, edit, or define implied foreign keys.

The Data Modeler toolbar icons are:

- **Save Diagram:** Saves the currently selected diagram. Diagrams are stored in a table that are created in the schema of the user.
- **Print Diagram:** Prints the selected diagram.
- **Save to SVG:** Saves the currently selected diagram to an image file in SVG format.
- **Add Objects to Diagram:** A dialog is displayed where you can select one or more objects from a specific schema into the selected diagram. Type * or a few characters in the search field to list the objects from the selected schema.
- **Add Note:** Adds notes to the selected diagram. Select the note to see the associated properties in the right pane, such as Text and Colors.
 - To enter text, expand the TEXT box and type information.
 - To select the background, border, or text color for the note, click the box and select the required color. To revert to the default color, click the Reset icon.
- **Delete:** Deletes the selected object or objects from the diagram. To select multiple objects, press the CTRL key and select the objects.
- **Zoom In:** Displays more detail, and potentially fewer objects, in the currently selected diagram.
- **Zoom Out:** Displays less detail, and potentially more objects, in the currently selected diagram.
- **Fit Screen:** Makes all relevant objects fit in the window for the currently selected diagram, adjusting the sizes of shapes and text labels as needed.
- **Actual Size:** Adjusts the shapes and text labels in the currently selected diagram to the default sizes.
- **Layout: Auto Layout:** Rearranges the objects in the diagram to a layout that may be more meaningful and attractive. If you do not like the rearrangement, you can restore the previous layout by clicking Undo.
- **Layout: Star Layout:** Rearranges the objects in the diagram to a star schema layout, where the fact table is in the center and the associated dimension tables surround the fact table.
- **DDL Preview:** Shows the DDL statements that would be generated to create the object. You can preview the DDL for selected objects by dragging the cursor over the objects that you want to preview and clicking **DDL Preview**. You have the option to save or send the DDL statements to the worksheet. To specify options on how to generate the DDL, click **Options**.
- **DDL Preview for Current Schema:** Shows the DDL statements that would be generated for the current schema. You can send the DDL statements to the worksheet. To specify DDL options for the schema, click **Options**.
- **Diagram Report:** Generates a report for tables and views in the selected diagram.
- **Schema Report:** Generates a report for tables and views in the selected schema.
- **Help:** Displays the help for the Data Modeling editor.

7

Creating and Editing Database Objects

You can create and edit objects for some object types by using the context menu in the Navigator tab. If you do not have any existing objects for the object type (such as tables), you can create objects by using the New Object icon, after first selecting the object type from the drop-down list. For all objects, the context menu includes **Quick DDL**, which saves the DDL statements for creating the object to a worksheet or to a text file. The various object types and the corresponding options are:

- **Tables:** Tables are used to hold data. Each table has multiple columns that describe the attributes of the database entity associated with the table, and each column has an associated data type.

You can create new tables, edit existing tables, and add indexes by using the options available in the context menu. Or you can select **Use as Template**, which displays the dialog box with information for the selected object filled in. You should then change the object name and any other aspects of the definition, as needed, to create a new object of that type. For a description of the fields, see [The Table Properties Dialog](#).

- **Views:** Views are virtual tables that select data from one or more underlying tables. For information about creating and editing a view, see [The View Properties Dialog](#).
- **Indexes:** An index contains an entry for each value that appears in the indexed column of the table and provides direct, fast access to rows. For information about creating and editing an index for a table, see [The Index Properties Dialog](#).
- **Packages, Functions, Procedures:** A package contains subprograms, which are programming objects that can be stored and executed in the database server, and called from other programming objects or applications. A package can contain functions or procedures, or both. Select **Run** from the context menu to insert the code for a selected function, procedure, or package into the worksheet.
- **Triggers:** Triggers are stored PL/SQL blocks associated with a table, a schema, the database, or anonymous PL/SQL blocks or calls to a procedure implemented in PL/SQL or Java.
- **Types:** A data type associates a fixed set of properties with the values that can be used in a column of a table or in an argument of a function or procedure.
- **Sequences:** Sequences are used to generate unique integers. You can use sequences to automatically generate primary key values. For information about creating and editing a sequence, see [The Sequence Properties Dialog](#).
- **Materialized View:** A materialized view is a database object that contains the results of a query. The FROM clause of the query can name tables, views, and other materialized views. For information about creating and editing a materialized view, see [Materialized View](#).
- **Materialized Views Logs:** A materialized view log is a table associated with the master table of a materialized view. When DML changes are made to master table data, Oracle Database stores rows describing those changes in the materialized

view log and then uses the materialized view log to refresh materialized views based on the master table.

- **Synonyms, Public Synonyms:** Synonyms provide alternative names for tables, views, sequences, procedures, stored functions, packages, materialized views, Java class database objects, user-defined object types, or other synonyms. For information about creating a synonym, see [The Synonym Properties Dialog](#).
- **Database Links:** A database link is a database object in one database that enables you to access objects on another database. After you have created a database link, you can use it to refer to tables and views in the other database.

7.1 The Table Properties Dialog

The Table Properties Dialog is displayed when you create a table, edit an existing table, or create a table using an existing one as a template. To create or edit a table, right-click a table object in the Navigator tab. The table options are also available when you right-click a table object in the Data Modeler Navigator tab. The table properties are grouped in several panes.

If you are editing an existing table, you can visit the panes in any order. If you click Create before you are finished creating the table, right-click the table name, select **Edit**, and continue creating the table.

Note:

Editing a partitioned table is not recommended. To identify whether a table is partitioned or not, right-click the table name and select **Edit**. If the table is partitioned, a warning message will be displayed.

Schema: Database schema in which to create the table. By default, a new table is created in the existing schema or the schema that you are logged into.

Name: Name for the table.

The different panes in the dialog are described in the following sections:

- [Columns pane](#)
- [Primary Key pane](#)
- [Unique Keys pane](#)
- [Indexes pane](#)
- [Foreign Keys pane](#)
- [Table Constraints pane](#)
- [Comments pane](#)
- [Storage pane](#)
- [External Table Properties pane](#)
- [Materialized View](#)
- [DDL pane](#)
- [Output pane](#)

Columns pane

Specifies properties for each column in the table.

General tab

Lists the columns currently in the table. To add a column, click the Add Column (+) icon, from the table below, select the column whose properties you want to add, and enter the details for each column in the fields at the bottom. To delete a column, select the row and click the Remove Column (-) icon; to move a column up or down in the table definition, select it and use the up-arrow and down-arrow buttons.

- **Name:** Name for the column.
- **Datatype:** Data type for the column.
- **Default:** If no value is specified, the default value inserted into the column when a row is inserted.
- **Default on NULL:** Applicable for Oracle Database 12c and later releases. If this option is selected, when a row is inserted into the table and the value specified for the column is NULL, the default value is inserted into the column.
- **Expression:** Expression for computing the value in the column.
- **Comments:** Optional descriptive comments about the column. Use this field to provide descriptions for the attributes.

In the table:

- **PK:** If this option is selected, the column becomes the primary key.
- **Identity Column:** If this option is selected, the column becomes an identity column. This is applicable only for Oracle Database 12c and later releases. For more details, see the Identity Column tab.

Constraints tab

Displays the Not Null and Check Constraints for a column. A check constraint requires values in a column to comply with a specified condition.

- **Not Null Constraint: Name:** Name for the Not Null constraint.
- **Not Null Constraint: Not Null:** If this option is selected, the column must contain data. You cannot specify no value or an explicit null value for this column when you insert a row. If this option is not checked, the column can contain either data or no data. A primary key column cannot be null.
- **Check Constraint: Name:** Name for the check constraint definition.
- **Check Constraint: Constraint:** Condition that must be met for a column to fulfill the check constraint. You can use any valid CHECK clause (without the CHECK keyword). For example, to indicate that the value in a numeric column named RATING must be from 1 to 10, you can specify: rating >=1 and rating <= 10.
- **Enabled:** If this option is selected, the constraint is checked when data is entered or updated in the column.
- **Deferrable:** If this option is selected, you can defer checking the validity of the constraint until the end of a transaction.
- **Initially Immediate:** If this option is selected, the constraint is checked whenever you add, update, or delete data from the column.

- **Validate:** If this option is selected, the existing data is checked to see if it conforms to the constraint.

Identity Column tab

Applicable for Oracle Database 12c and later releases. The Identity Column tab lists the properties of the identity column. This tab becomes available only after the Identity Column checkbox is selected for the column in the General tab. An identity column is an autoincrement column that can be used to identify a table row. Only one identity column can be specified for a table.

- **Generate:** Always means that values cannot be explicitly included for the identity column in INSERT OR UPDATE statements, By Default means values for the identity column are generated automatically if no values are specified explicitly, By Default on Null means values are generated for the column only when a NULL value is supplied.
- **Start with:** Starting value of the sequence.
- **Increment:** Interval between successive numbers in a sequence.
- **Min value:** Lowest possible value for the sequence. The default is 1 for an ascending sequence and $-(10^{26})$ for a descending sequence.
- **Max value:** Highest possible value for the sequence. The default is 10^{27} for an ascending sequence and -1 for a descending sequence.
- **Cache and Cache size:** Cache causes sequence values to be preallocated in cache, which can improve application performance; Cache size indicates the number of sequence values preallocated in cache. No Cache causes sequence values not to be preallocated in cache.
- **Cycle:** Indicates whether the sequence "wraps around" to reuse numbers after reaching its maximum value (for an ascending sequence) or its minimum value (for a descending sequence). If cycling of values is not enabled, the sequence cannot generate more values after reaching its maximum or minimum value.
- **Order:** Indicates whether sequence numbers are generated in the order in which they are requested. If No Order is specified, sequence numbers are not guaranteed to be in the order in which they were requested.

Primary Key pane

Specifies the primary key for the table. The primary key is the column, or set of columns, that uniquely identifies each row in the table. If the Primary Key checkbox is selected for a column in the General tab, the corresponding fields are automatically populated in the Primary Key pane. You can make changes to the properties as required.

An index is automatically created on the primary key.

- **Name:** Name of the constraint to be associated with the primary key definition.
- **Enabled:** If this option is checked, the primary key constraint is enforced: that is, the data in the primary key column (or set of columns) must be unique and not null.
- **Index:** Name of the index to which the primary key refers.
- **Tablespace:** Name of the tablespace associated with the index.

- **Available Columns:** Lists the columns that are available to be added to the primary key definition. You can select multiple attributes, if required, for the primary key.
- **Selected Columns:** Lists the columns that are included in the primary key definition.

To add a column to the primary key definition, select it in Available Columns and click the Add (>) icon; to remove a column from the primary key definition, select it in Selected Columns and click the Remove (<) icon. To move all columns from available to selected (or the reverse), use the Add All (>>) or Remove All (<<) icon. To move a column up or down in the primary key definition, select it in Selected Columns and use the arrow buttons.

Unique Keys pane

Specifies one or more unique constraints for the table. A unique constraint specifies a column, or set of columns, whose data values must be unique: each data value must not be null, and it must not be the same as any other value in the column.

To add a unique constraint, click the Add button; to delete a unique constraint, select it and click the Remove button.

- **Name:** Name of the unique constraint.
- **Enabled:** If this option is selected, the unique constraint is enforced.
- **Rely:** If this option is selected, the constraint in NOVALIDATE mode is taken into account during query rewrite.
- **Deferrable:** If this option is selected, in subsequent transactions, constraint checking can be deferred until the end of the transaction using the SET CONSTRAINT(S) statement.
- **Initially Immediate:** If this option is selected, the constraint is checked at the end of each subsequent SQL statement.
- **Validate:** If the option is selected, the existing data is checked to see if it conforms to the constraint.
- **Index:** Name of the index to which the unique key refers.
- **Tablespace:** Name of the tablespace associated with the index.
- **Available Columns:** Lists the columns that are available to be added to the unique constraint definition.
- **Selected Columns:** Lists the columns that are included in the unique constraint definition.

To add a column to the unique constraint definition, select it in Available Columns and click the Add (>) icon; to remove a column from the unique constraint definition, select it in Selected Columns and click the Remove (<) icon. To move all columns from available to selected (or the reverse), use the Add All (>>) or Remove All (<<) icon. To move a column up or down in the unique constraint definition, select it in Selected Columns and use the arrow buttons.

Indexes pane

Lists the indexes currently defined on the table. To add an index, click the Add Index (+) icon; to delete an index, select it and click the Remove Index (-) icon.

- **Name:** Name of the index.

- **Type:** The type of Oracle index. `Non-unique` means that the index can contain multiple identical values; `Unique` means that no duplicate values are permitted; `Bitmap` stores rowids associated with a key value as a bitmap.
- **Tablespace:** Name of the tablespace for the index.
- **Expression:** A column expression is an expression built from columns, constants, SQL functions, and user-defined functions. When you specify a column expression, you create a function-based index.
- **Available Columns** and **Selected Columns:** Columns selected for the index. To select a column, click the column in the Available Columns box, and then click the Add Selected Columns icon to move it to the Selected Columns box.

Foreign Keys pane

Specifies one or more foreign keys for the table. A foreign key specifies a column ("local column"), each of whose data values must match a value in the primary key or unique constraint of another table.

- **Name:** Name of the foreign key definition.
- **Enabled:** If this option is checked, the foreign key is enforced.
- **Rely, Deferrable, Initially Immediate, Validate:** See the description of these fields in the Unique Keys pane.
- **Referenced Constraint: Schema:** Name of the schema containing the table with the primary key or unique constraint to which this foreign key refers.
- **Referenced Constraint: Table:** Name of the table with the primary key or unique constraint to which this foreign key refers.
- **Referenced Constraint: Constraint:** Name of the primary key or unique constraint to which this foreign key refers.
- **Referenced Constraint: On Delete:** Action to take automatically when a row in the referenced table is deleted and rows with that value exist in the table containing this foreign key: `NO ACTION` (shown by a crossing line in diagrams) performs no action on these rows; `CASCADE` (shown by an "X") deletes these rows; `SET NULL` (shown by a small circle) sets null all columns in those rows that can be set to a null value.
- **Associations: Local Column:** Lists the column in the currently selected (local) table that is included in the foreign key definition. For each referenced column in the foreign key definition, select the name of a column in the edited table.
- **Associations: Referenced Column:** For each local column, identifies the column in the other (foreign) table that must have a value matching the value in the local column.

Table Constraints pane

Specifies one or more check constraints for the table. A check constraint specifies a condition that must be met when a row is inserted into the table or when an existing row is modified.

- **Name:** Name of the check constraint definition.
- **Check Condition:** Condition that must be met for a row to fulfil the check constraint. You can use any valid CHECK clause (without the CHECK keyword).

For example, to indicate that the value in a numeric column named RATING must be from 1 to 10, you can specify rating >=1 and rating <= 10.

- **Enabled:** If this option is checked, the check constraint is enforced.

Comments pane

Optional descriptive comments about the table.

Storage pane

Enables you to specify storage options for the table. When you create or edit a table or an index, you can override the default storage options.

- **Organization:** Specifies that the table is stored and organized with (Index) or without an index (Heap) or as an external table (External).
- **Tablespace:** Name of the tablespace for the table or index.
- **Logging:** ON means that the table creation and any subsequent INSERT operations against the table are logged in the redo log file. OFF means that these operations are not logged in the redo log file.
- **Row Archival:** YES enables in-database archiving, which allows you to archive rows within the table by marking them as invisible.

External Table Properties pane

Specifies options for an external table, which is a read-only table whose metadata is stored in the database but whose data is stored outside the database.

External Table

- **Access Driver Type:** Specifies the type of external table.
 - ORACLE_LOADER: Extracts data from text data files. This is the default access driver, which loads data from external tables to internal tables.
 - ORACLE_DATAPUMP: Extracts data from binary dump files. This access driver can perform both loads and unloads.
 - ORACLE_BIGDATA: Extracts data from Oracle Big Data Appliance.
 - ORACLE_HDFS: Extracts data stored in a Hadoop Distributed File System (HDFS).
 - ORACLE_HIVE: Extracts data stored in Apache HIVE.
- **Default Directory:** Specifies the default directory to use for all input and output files that do not explicitly name a directory object. The location is specified with a directory object, not a directory path.
- **Access Params:** Assigns values to the parameters of the specific access driver for the external table. Access parameters are optional.
 - OPAQUE_FORMAT_SPEC: The opaque_format_spec specifies all access parameters for the ORACLE_LOADER, ORACLE_DATAPUMP, ORACLE_HDFS, and ORACLE_HIVE access drivers. For descriptions of the access parameters, see *Oracle Database Utilities*. Field names specified in the opaque_format_spec must match columns in the table definition, else Oracle Database ignores them.

- USING CLOB: Enables you to derive the parameters and their values through a subquery. The subquery cannot contain any set operators or an ORDER BY clause. It must return one row containing a single item of data type CLOB.
- **Reject Limit:** The number of conversion errors that can occur during a query of the external data before an Oracle Database error is returned and the query is aborted.
- **Project Column:** Determines how the access driver validates the rows of an external table in subsequent queries.
 - ALL: Processes all column values, regardless of which columns are selected, and validates only those rows with fully valid column entries. If any column value raises an error, such as a data type conversion error, the row is rejected even if that column was not referenced in the select list of the query.
 - REFERENCED: Processes only those columns in the select list of the query.

The ALL setting guarantees consistent result sets. The REFERENCED setting can result in different numbers of rows returned, depending on the columns referenced in subsequent queries, but is faster than the ALL setting. If a subsequent query selects all columns of the external table, then the settings behave identically.

- **Location:** Specifies the data files for the external table. Use the Add (+) icon to add each location specification.
 - For ORACLE_LOADER and ORACLE_DATAPUMP, the files are named in the form `directory:file`. The directory portion is optional. If it is missing, then the default directory is used as the directory for the file. If you are using the ORACLE_LOADER access driver, then you can use wildcards in the file name. An asterisk (*) signifies multiple characters and a question mark (?) signifies a single character.
 - For ORACLE_HDFS, LOCATION is a list of Uniform Resource Identifiers (URIs) for a directory or for a file. There is no directory object associated with a URI.
 - For ORACLE_HIVE, LOCATION is not used. Instead, the Hadoop HCatalog table is read to obtain information about the location of the data source (which could be a file or another database).

Opaque Format Spec

Specifies all access parameters for the ORACLE_LOADER, ORACLE_DATAPUMP, ORACLE_HDFS, and ORACLE_HIVE access drivers.

For example:

```
RECORDS DELIMITED BY NEWLINE CHARACTERSET US7ASCII
  TERRITORY AMERICAN
  BADFILE log_file_dir:'ext_lv3.bad'
  LOGFILE log_file_dir:'ext_lv3.log'
  FIELDS TERMINATED BY "|" OPTIONALLY ENCLOSED BY '^' LDRTRIM
  ( PROD_ID,
    CUST_ID ,
    TIME_ID DATE(10) "YYYY-MM-DD",
    CHANNEL_ID ,
    PROMO_ID ,
    QUANTITY_SOLD ,
    AMOUNT_SOLD ,
```

```
        UNIT_COST ,  
        UNIT_PRICE  
    )
```

and the full statement:

```
CREATE TABLE SH.SALES_TRANSACTIONS_EXT  
(  
    PROD_ID NUMBER ,  
    CUST_ID NUMBER ,  
    TIME_ID DATE ,  
    CHANNEL_ID NUMBER ,  
    PROMO_ID NUMBER ,  
    QUANTITY_SOLD NUMBER ,  
    AMOUNT_SOLD NUMBER (10,2) ,  
    UNIT_COST NUMBER (10,2) ,  
    UNIT_PRICE NUMBER (10,2)  
)  
ORGANIZATION EXTERNAL  
(  
    TYPE ORACLE_LOADER  
    DEFAULT DIRECTORY DATA_FILE_DIR  
    ACCESS PARAMETERS  
(  
    RECORDS DELIMITED BY NEWLINE CHARACTERSET US7ASCII  
    TERRITORY AMERICAN  
    BADFILE log_file_dir:'ext_lv3.bad'  
    LOGFILE log_file_dir:'ext_lv3.log'  
    FIELDS TERMINATED BY "|" OPTIONALLY ENCLOSED BY '^' LDRTRIM  
( PROD_ID ,  
    CUST_ID ,  
    TIME_ID DATE(10) "YYYY-MM-DD",  
    CHANNEL_ID ,  
    PROMO_ID ,  
    QUANTITY_SOLD ,  
    AMOUNT_SOLD ,  
    UNIT_COST ,  
    UNIT_PRICE  
)  
)  
LOCATION ( "DATA_FILE_DIR":'salelv3.dat')  
)  
REJECT LIMIT 100  
;
```

CLOB Subquery

Type or copy and paste the query.

 **Note:**

For more information about the external table fields, see *Database Utilities* and *SQL Language Reference*.

Materialized View

Query: Contains the SQL code for the query part of the view definition. Type or copy and paste the query.

General

- **On Pre-built Table:** If **Yes**, an existing table is registered as a preinitialized materialized view. This option is particularly useful for registering large materialized views in a data warehousing environment. The table must have the same name and be in the same schema as the resulting materialized view, and the table should reflect the materialization of a subquery.
- **Reduced Precision:** **Yes** authorizes the loss of precision that will result if the precision of the table or materialized view columns do not exactly match the precision returned by the subquery. If **No**, the precision of the table or materialized view columns must exactly match the precision returned by the subquery, or the create operation will fail.
- **For Update:** Select **Yes** to allow a subquery, primary key, object, or rowid materialized view to be updated. When used in conjunction with Advanced Replication, these updates will be propagated to the master.
- **Real Time MV:** Select **Yes** to create a real-time materialized view or a regular view. A real-time materialized view provides fresh data to user queries even when the materialized view is not in sync with its base tables due to data changes. Instead of modifying the materialized view, the optimizer writes a query that combines the existing rows in the materialized view with changes recorded in log files (either materialized view logs or the direct loader logs). This is called on-query computation.
- **Query Rewrite:** If **Enable**, the materialized view is enabled for query rewrite, which transforms a user request written in terms of master tables into a semantically equivalent request that includes one or more materialized views.
- **Build:** Specifies when to populate the materialized view. **Immediate** indicates that the materialized view is to be populated immediately. **Deferred** indicates that the materialized view is to be populated by the next refresh operation. If you specify **Deferred**, the first (deferred) refresh must always be a complete refresh; until then, the materialized view has a staleness value of unusable, so it cannot be used for query rewrite.
- **Use Index:** If **Yes**, a default index is created and used to speed up incremental (fast) refresh of the materialized view. If **No**, this default index is not created. (For example, you might choose to suppress the index creation now and to create such an index explicitly later.)
- **Index Tablespace:** Specifies the tablespace in which the materialized view is to be created. If a tablespace is not selected, the materialized view is created in the default tablespace of the schema containing the materialized view.
- **Cache:** If **Yes**, the blocks retrieved for this table are placed at the most recently used end of the least recently used (LRU) list in the buffer cache when a full table

scan is performed. This setting is useful for small lookup tables. If **No**, the blocks are placed at the least recently used end of the LRU list.

Refresh Clause

- **Refresh:** Select Yes to enable refresh operations.
- **Refresh Type:** The method of refresh operation to be performed:
 - Complete Refresh: Executes the defining query of the materialized view, even if a fast refresh is possible.
 - Fast Refresh: Uses the incremental refresh method, which performs the refresh according to the changes that have occurred to the master tables. The changes for conventional DML changes are stored in the materialized view log associated with the master table. The changes for direct-path INSERT operations are stored in the direct loader log.
 - Force Refresh: Performs a fast refresh if one is possible; otherwise, performs a complete refresh.
- **Action:** The type of refresh operation to be performed:
 - On Demand: Performs a refresh when one of the DBMS_MVIEW refresh procedures are called.
 - On Commit: Performs a fast refresh whenever the database commits a transaction that operates on a master table of the materialized view. This may increase the time taken to complete the commit, because the database performs the refresh operation as part of the commit process.
 - Specify: Performs refresh operations according to what you specify in the Start on and Next fields.
- **Start Date:** Starting date and time for the first automatic refresh operation. Must be in the future.
- **Next Date:** Time for the next automatic refresh operation. The interval between the Start on and Next times establishes the interval for subsequent automatic refresh operations. If you do not specify a value, the refresh operation is performed only once at the time specified for Start on.
- **With:** Refresh type, which determines the type of materialized view:
 - Primary Key: Creates a primary key materialized view, which allows materialized view master tables to be reorganized without affecting the eligibility of the materialized view for fast refresh.
 - Row ID: Creates a rowid materialized view, which is useful if the materialized view does not include all primary key columns of the master tables.
- **Default Storage:** If Yes, DEFAULT specifies that Oracle Database will choose automatically which rollback segment to use. If you specify DEFAULT, you cannot specify the rollback_segment. DEFAULT is most useful when modifying, rather than creating, a materialized view.
- **Storage Type:** MASTER specifies the remote rollback segment to be used at the remote master site for the individual materialized view. LOCAL specifies the remote rollback segment to be used for the local refresh group that contains the materialized view. This is the default.
- **Rollback Segment:** Enter the name of the rollback segment.

- **Using Constraint:** If this option is checked, more rewrite alternatives can be used during the refresh operation, resulting in more efficient refresh execution. The behavior of this option is affected by whether you select Enforced or Trusted.
 - **Enforced:** Causes only enforced constraints to be used during the refresh operation.
 - **Trusted:** Enables the use of dimension and constraint information that has been declared trustworthy by the database administrator but that has not been validated by the database. If the dimension and constraint information is valid, performance may improve. However, if this information is invalid, then the refresh procedure may corrupt the materialized view even though it returns a success status.

DDL pane

You can review and save the SQL statements that SQL Developer Web will use to create or edit the object. If you want to make any changes, go back to the relevant panes and make the changes there.

For a new table, click **CREATE** to view the generated DDL statements. When you edit table properties, click **UPDATE** to view the generated ALTER statements. For a new table, the UPDATE tab will not be available.

When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the appropriate pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

Related Topics

- *Oracle Database SQL Language Reference*

7.2 The Index Properties Dialog

The Index Properties dialog box is displayed when you create or edit an index. You can create and edit in three ways:

- Right-click a table in the worksheet Navigator tab, and select **Add Index**. To edit, right-click an index in the Navigator tab and select **Edit**.
- Right-click a table object in the Data Modeler Navigator tab, and select **Add Index**.
- Create and edit an index in the Table Properties Dialog (See [The Table Properties Dialog](#)).

Definition pane

- **Schema:** Database schema that owns the table associated with the index.
- **Table:** Name of the table associated with the index.
- **Schema:** Database in which to create the index.
- **Tablespace:** Tablespace for the index.
- **Name:** Name of the index.

- **Type:** The type of Oracle index. **Non-unique** means that the index can contain multiple identical values; **Unique** means that no duplicate values are permitted; **Bitmap** stores rowids associated with a key value as a bitmap.
- **Expression:** A column name or column expression. A column expression is an expression built from columns, constants, SQL functions, and user-defined functions. When you specify a column expression, you create a function-based index.
- **Available Columns:** Columns available in the table.
- **Selected Columns:** Columns selected for the index. Click the Add Selected Columns icon to move columns from the Available Columns list.
- **Order:** **ASC** for an ascending index (index values sorted in ascending order); **DESC** for a descending index (index values sorted in descending order).

DDL pane

You can review and save the SQL statements that SQL Developer Web will use to create or edit the index. If you want to make any changes, go back to the Definition pane and make the changes there. For a new index, click **CREATE** to view the generated DDL statements. When you edit index properties, click **UPDATE** to view the generated ALTER statements. For a new index, the UPDATE tab will not be available. When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the Definition pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

7.3 The Sequence Properties Dialog

This Sequence Properties Dialog is displayed when you right-click a sequence object in the worksheet Navigator tab and select **Add Sequence** or **Edit Sequence**. These options are also available when you right-click a table object in the Data Modeler Navigator tab.

A sequence is an object from which multiple users may generate unique integers. You can use sequences to automatically generate primary key values.

Properties pane

- **Schema:** Database schema in which to create the sequence.
- **Name:** Name of the sequence.
- **Start with:** Starting value of the sequence.
- **Increment:** Interval between successive numbers in a sequence.
- **Min value:** Lowest possible value for the sequence. The default is 1 for an ascending sequence and $-(10^{26})$ for a descending sequence.
- **Max value:** Highest possible value for the sequence. The default is 10^{27} for an ascending sequence and -1 for a descending sequence.
- **Cache** and **Cache size:** **Cache** causes sequence values to be preallocated in cache, which can improve application performance; **Cache size** indicates the number of sequence values preallocated in cache. **No Cache** causes sequence values not to be preallocated in cache.

- **Cycle:** Indicates whether the sequence "wraps around" to reuse numbers after reaching its maximum value (for an ascending sequence) or its minimum value (for a descending sequence). If cycling of values is not enabled, the sequence cannot generate more values after reaching its maximum or minimum value.
- **Order:** Indicates whether sequence numbers are generated in the order in which they are requested. If **No Order** is specified, sequence numbers are not guaranteed to be in the order in which they were requested.

DDL pane

You can review and save the SQL statements that SQL Developer Web will use to create or edit the sequence. If you want to make any changes, go back to the Properties pane and make the changes there. For a new sequence, click **CREATE** to view the generated DDL statements. When you edit a sequence, click **UPDATE** to view the generated ALTER statements. For a new sequence, the UPDATE tab will not be available. When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the Properties pane, fix the errors, and run the commands again. You have save to a text file or clear the output.

7.4 The View Properties Dialog

The View Properties Dialog is displayed when you create or edit a view. To create or edit, right-click a view in the worksheet Navigator tab, and select **New** or **Edit**.

Schema: Database schema in which to create the view.

Name: Name of the view.

The different panes in the dialog are described in the following sections:

SQL Query pane

Enter or copy and paste the SQL query for the view, using the SELECT and FROM keywords along with the syntax needed to retrieve the desired information. A semicolon is not required after the query.

Columns pane

Click the Refresh Columns icon to automatically populate the columns in this pane. You can edit the columns by selecting the required row and making changes in the Header Alias and Comments fields.

Storage pane

- **Force on Create:** Select **Yes** to create the view regardless of whether the base tables of the view or the referenced object types exist or the owner of the schema containing the view has privileges on them. These conditions must be true before any SELECT, INSERT, UPDATE, or DELETE statements can be issued against the view. If the view definition contains any constraints, CREATE VIEW ... FORCE fails if the base table does not exist or the referenced object type does not exist. CREATE VIEW ... FORCE also fails if the view definition names a constraint that does not exist.

- **Query Restriction: Read Only** prevents the view from being used to add, delete, or change data in the underlying table. **Check Option** prohibits any changes to the underlying table that would produce rows that are not included in this view.

Use the Primary Key, Unique Keys, Foreign Keys, and Comments panes to add or edit properties as required.

DDL pane

Based on the inputs provided, the DDL statements are generated. You can review and save the SQL statements. If you want to make any changes, go back to the relevant pane and make the changes there. For a new view, click **CREATE** to view the generated DDL statements. When you edit a view, click **UPDATE** to view the generated ALTER statements. For a new view, the UPDATE tab will not be available. When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the respective pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

7.5 The Synonym Properties Dialog

The Synonym Properties Dialog is displayed when you create a synonym.

There are two ways of creating a synonym:

- In the worksheet navigator pane, right-click the object for which you want to create the synonym, and select **Add Synonym**. In this case, the only fields that you can edit in the Properties pane are Public and Synonym Name. The values of the remaining fields are predetermined by the object selected.
- In the worksheet navigator pane, select the object type as **Synonyms** or **Public Synonyms** from the drop-down list. Click the New Object (+) icon. All the fields in the Properties dialog are available for edit.

The different panes in the dialog are described in the following sections:

Properties pane

- **Public:** If this option is checked, the synonym is accessible to all users. However, each user must have appropriate privileges on the underlying object to use the synonym. If this option is not checked, the synonym is a private synonym, and is accessible only within its schema.
- **Synonym Schema:** Database schema in which to create the synonym.
- **Synonym Name:** Name of the synonym. A private synonym must be unique within its schema; a public synonym must be unique within the database.
- **Object Type:** Specify the type of object to which this synonym refers.
- **Object Schema:** Schema containing the object or name to which this synonym refers.
- **DB Filter:** After selecting the Object Type and Object Schema, the list of objects of the selected type may be very long. To filter the object names, enter the search entry and click the Refresh icon. The Object Name field is auto-filled with appropriate object names in the drop-down list.

- **Object Name:** Select the name of the object to which this synonym refers.
- **DB Link:** Enter a complete or partial database link to create a synonym for a schema object on a remote database where the object is located. If you specify DB Link and omit schema, then the synonym refers to an object in the schema specified by the database link. Oracle recommends that you specify the schema containing the object in the remote database. If you omit DB Link, then Oracle Database assumes the object is located on the local database.

DDL pane

Based on the inputs provided, the DDL statements are generated. You can review and save the SQL statements. If you want to make any changes, go back to the relevant pane and make the changes there. For a new view, click **CREATE** to view the generated DDL statements. When you edit a view, click **UPDATE** to view the generated ALTER statements. For a new view, the UPDATE tab will not be available. When you are finished, click **Apply**.

Output pane

Displays the results of the DDL commands. If there are any errors, go to the respective pane, fix the errors, and run the commands again. You can save to a text file or clear the output.

7.6 Implied Foreign Keys

Implied foreign keys are dependencies that exist between tables but are not defined in the database. In a data warehouse environment, it is a common practice not to create foreign keys. However, it becomes necessary to show these dependencies for presentation or reporting purposes.

In SQL Developer Web, you can display implied foreign keys for objects in a star schema by defining them or by discovering them in the data dictionary. It is possible to have more than one source for implied foreign keys.

Define Implied Foreign Keys

You can defined implied foreign keys in two ways:

- **Using the Implied Foreign Keys dialog**

You can define implied foreign keys using the Implied Foreign Keys dialog.

1. Right-click an object (table or view) in a diagram and select **Implied Foreign Keys**. The Implied Foreign Keys dialog is displayed.
2. In the Implied Foreign Keys dialog, click the + icon to add an entry in the grid.
3. Select the entry in the grid to enable and enter values in the following fields:
 - **Referenced Object:** Object in the diagram that has a dependency to the source object.
 - **Local Column:** Name of the column in the source object.
 - **Referenced Column:** Name of the column in the targeted object.
 - **Discovery Sources:** Automatically prefilled, displays whether the implied foreign keys have been defined or were discovered in the data dictionary.

4. Click **OK**. The implied foreign key dependency is displayed with a dotted line on the diagram.
- **By dragging the arrow to the referenced object in the diagram**
You can also define an implied foreign key in the following way:
 1. Select the source object on the diagram.
 2. Click and drag the small blue curved arrow to the referenced object. The dependency will be displayed with a dotted line on the diagram.
 3. Right-click the source object and enter the column names in the Implied Foreign Keys dialog.

Discover Implied Foreign Keys in the Data Dictionary

SQL Developer Web can automatically create a star schema by searching for several types of definitions in the data dictionary.

In the Navigator tab, right-click an object (table or view) and select **Add Object as Star Schema to Diagram**. The object must be a fact table. The data dictionary is then searched for joins and dependencies related to the object, such as:

- Foreign keys defined for the selected table to other tables. If implied foreign keys are later discovered for the same columns, they will not be displayed on the diagram.
- Joins used in the definitions of bitmap join indexes, materialized views with aggregates, and attribute clustering.
- Dependencies based on dimension definitions and column name matching in fact table.
- Fact and dimension definitions for Analytic views and OLAP cube and dimensions.