

Oracle® Database

Graph Developer's Guide for Property Graph



21.4
F47626-04
May 2022

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

ORACLE®

Oracle Database Graph Developer's Guide for Property Graph, 21.4

F47626-04

Copyright © 2016, 2022, Oracle and/or its affiliates.

Primary Author: Lavanya Jayapalan

Contributors: Prashant Kannan, Chuck Murray, Melliyal Annamalai, Korbinian Schmid, Albert Godfrind, Oskar van Rest, Jorge Barba, Ana Estrada, Steve Serra, Ryota Yamanaka, Bill Beauregard, Hector Briseno, Hassan Chafi, Eugene Chong, Souripriya Das, Juan Garcia, Florian Gratzer, Zazhil Herena, Sungpack Hong, Roberto Infante, Hugo Labra, Gabriela Montiel-Moreno, Eduardo Pacheco, Joao Paiva, Matthew Perry, Diego Ramirez, Siva Ravada, Carlos Reyes, Jane Tao, Edgar Vazquez, Zhe (Alan) Wu

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

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

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

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

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

Contents

Preface

Audience	XX
Documentation Accessibility	XX
Related Documents	XX
Conventions	XX

Changes in This Release for This Guide

Part I Getting Started with Oracle Property Graphs

1 Property Graph Support Overview

1.1 Introduction to Property Graphs	1-2
1.1.1 What Are Property Graphs?	1-2
1.1.2 About the Property Graph Feature of Oracle Database	1-3
1.1.2.1 In-Memory Graph Server (PGX)	1-4
1.1.2.2 Options for Property Graph Architecture	1-4
1.2 Before You Begin with Oracle Property Graphs	1-7
1.2.1 Database Compatibility and Restrictions	1-7
1.2.2 Downloading Oracle Graph Server and Client	1-8
1.2.3 Installing PL/SQL Packages in Oracle Database	1-9
1.3 Oracle Graph Server and Client Installation	1-10
1.3.1 Installing Oracle Graph Server	1-10
1.3.2 Deploying Oracle Graph Server to a Web Server	1-13
1.3.2.1 Deploying to Apache Tomcat	1-13
1.3.2.2 Deploying to Oracle WebLogic Server	1-14
1.3.3 Upgrading Oracle Graph Server	1-15
1.3.4 Uninstalling Oracle Graph Server	1-15
1.3.5 Installing Oracle Graph Client	1-15
1.3.5.1 Installing the Java Client	1-16
1.3.5.2 Installing the Python Client	1-17

1.3.5.3	Uninstalling the Python Client	1-18
1.3.5.4	Enabling the Graph Visualization Application	1-19
1.3.5.5	Deploying the Graph Visualization Application	1-19
1.3.5.6	Installing the Graph Zeppelin Interpreter Client	1-24
1.4	Setting Up Transport Layer Security	1-25
1.4.1	Generating a Self-Signed Server Certificate	1-25
1.4.2	Configuring the Graph Server (PGX)	1-26
1.4.3	Configuring a Client to Trust the Self-Signed Certificate	1-27
1.5	Adding Permissions to Publish the Graph	1-28
1.6	Security Best Practices with Graph Data	1-29
1.7	Interactive Graph Shell CLIs	1-29
1.7.1	Starting the OPG4J Shell	1-30
1.7.2	Starting the OPG4Py Shell	1-32
1.8	Developing Applications Using Graph Server Functionality as a Library	1-33
1.9	Storing Graphs in Oracle Database and Loading Graphs into Memory	1-34
1.9.1	Two-Tier Mode	1-34
1.9.2	Three-Tier Mode	1-34
1.10	Using Oracle Graph with the Autonomous Database	1-35
1.10.1	Two-Tier Deployments of Oracle Graph with Autonomous Database	1-36
1.10.2	Three-Tier Deployments of Oracle Graph with Autonomous Database	1-37
1.11	Migrating Property Graph Applications from Before Release 21c	1-39
1.12	Upgrading From Graph Server and Client 20.4.x to 21.x	1-42
1.13	Using the Graph Zeppelin Interpreter Client	1-44
1.14	About Oracle Graph Server and Client Accessibility	1-45

2 Quick Starts for Using Oracle Property Graph

2.1	Quick Start: Interactively Analyze Graph Data	2-1
2.1.1	Quick Start: Create and Query a Graph in the Database, Load into In-Memory Graph Server (PGX) for Analytics	2-1
2.1.1.1	Create and Query a Graph in the Database	2-2
2.1.1.2	Load the Graph into Memory and Run Graph Analytics	2-6
2.1.2	Quick Start: Create, Query, and Analyze a Graph in In-Memory Graph Server (PGX)	2-9
2.1.3	Quick Start: Executing PGQL Queries in SQLcl	2-14
2.2	QuickStart: Run Graph Analytics Using the Python Shell	2-14
2.3	Quick Start: Using the Python Client as a Module	2-15
2.4	Oracle LiveLabs Workshops for Graphs	2-17

3 Property Graph Views on Oracle Database Tables

3.1	Loading a Graph into the Graph Server (PGX) from a Property Graph View	3-4
-----	--	-----

4 Using the In-Memory Graph Server (PGX)

4.1	Overview of the In-Memory Graph Server (PGX)	4-2
4.1.1	Design of the In-Memory Graph Server (PGX)	4-2
4.1.2	Usage Modes of the In-memory Graph Server (PGX)	4-4
4.2	User Authentication and Authorization	4-6
4.2.1	Privileges and Roles in Oracle Database	4-7
4.2.2	Basic Steps for Using an Oracle Database for Authentication	4-7
4.2.3	Prepare the Graph Server for Database Authentication	4-9
4.2.4	Store the Database Password in a Keystore	4-11
4.2.5	Token Expiration	4-16
4.2.6	Advanced Access Configuration	4-17
4.2.7	Customizing Roles and Permissions	4-18
4.2.7.1	Checking Graph Permissions Using API	4-18
4.2.7.2	Adding and Removing Roles	4-21
4.2.7.3	Defining Permissions for Individual Users	4-21
4.2.7.4	Defining Permissions to Use Custom Graph Algorithms	4-21
4.2.8	Revoking Access to the Graph Server	4-21
4.2.9	Examples of Custom Authorization Rules	4-22
4.2.10	Kerberos Enabled Authentication	4-24
4.2.10.1	Prerequisite Requirements	4-24
4.2.10.2	Prepare the Graph Server for Kerberos Authentication	4-24
4.2.10.3	Login to the Graph Server Using Kerberos Ticket	4-25
4.3	About Vertex and Edge IDs	4-27
4.4	Reading Graphs from Oracle Database into the Graph Server (PGX)	4-29
4.4.1	Creating a JSON Configuration to Load a Graph	4-30
4.4.2	Building Graph Configuration Using the GraphConfigBuilder API	4-32
4.5	Keeping the Graph in Oracle Database Synchronized with the Graph Server	4-36
4.5.1	Examples of Synchronizing	4-37
4.6	Optimizing Graphs for Read Versus Updates in the In-Memory Graph Server (PGX)	4-42
4.7	Storing a Graph Snapshot on Disk	4-42
4.8	Executing Built-in Algorithms	4-46
4.8.1	About Built-In Algorithms in the In-Memory Graph Server (PGX)	4-47
4.8.2	Running the Triangle Counting Algorithm	4-47
4.8.3	Running the PageRank Algorithm	4-48
4.9	Using Custom PGX Graph Algorithms	4-49
4.9.1	Writing a Custom PGX Algorithm	4-49
4.9.1.1	Collections	4-50
4.9.1.2	Iteration	4-51

4.9.1.3	Reductions	4-51
4.9.2	Compiling and Running a Custom PGX Algorithm	4-52
4.9.3	Example Custom PGX Algorithm: PageRank	4-54
4.10	Creating Subgraphs	4-55
4.10.1	About Filter Expressions	4-55
4.10.2	Using a Simple Filter to Create a Subgraph	4-56
4.10.3	Using a Complex Filter to Create a Subgraph	4-57
4.10.4	Using a Vertex Set to Create a Bipartite Subgraph	4-58
4.11	Using Automatic Delta Refresh to Handle Database Changes	4-60
4.11.1	Configuring the In-Memory Server for Auto-Refresh	4-60
4.11.2	Configuring Basic Auto-Refresh	4-61
4.11.3	Reading the Graph Using the In-Memory Graph Server (PGX) or a Java Application	4-62
4.11.4	Checking Out a Specific Snapshot of the Graph	4-62
4.11.5	Advanced Auto-Refresh Configuration	4-63
4.11.6	Special Considerations When Using Auto-Refresh	4-64
4.12	Starting the In-Memory Graph Server (PGX)	4-64
4.12.1	Starting and Stopping the Graph Server (PGX) Using the Command Line	4-65
4.12.2	Configuring the In-Memory Graph Server (PGX)	4-65
4.13	Connecting to the In-Memory Graph Server (PGX)	4-72
4.13.1	Connecting with the Graph Shell	4-73
4.13.2	Connecting with Java	4-78
4.13.2.1	Starting and Stopping the PGX Engine	4-78
4.13.3	Connecting with Python	4-79
4.14	Using Graph Server (PGX) as a Library	4-80
4.15	User-Defined Functions (UDFs) in PGX	4-81
4.16	Using HAProxy for PGX Load Balancing and High Availability	4-85

5 Using the Property Graph Schema

5.1	Property Graph Schema Objects for Oracle Database	5-2
5.1.1	Property Graph Tables (Detailed Information)	5-2
5.1.2	Default Indexes on Vertex (VT\$) and Edge (GE\$) Tables	5-6
5.1.3	Flexibility in the Property Graph Schema	5-7
5.2	Data Access Layer	5-7
5.3	Getting Started with Property Graphs	5-7
5.3.1	Required Privileges for Database Users	5-8
5.4	Using Java APIs for Property Graph Data	5-8
5.4.1	Overview of the Java APIs	5-9
5.4.1.1	Oracle Graph Property Graph Java APIs	5-9
5.4.1.2	Oracle Database Property Graph Java APIs	5-9
5.4.2	Parallel Retrieval of Graph Data	5-9

5.4.3	Using an Element Filter Callback for Subgraph Extraction	5-11
5.4.4	Using Optimization Flags on Reads over Property Graph Data	5-14
5.4.5	Adding and Removing Attributes of a Property Graph Subgraph	5-16
5.4.6	Getting Property Graph Metadata	5-21
5.4.7	Merging New Data into an Existing Property Graph	5-22
5.4.8	Opening and Closing a Property Graph Instance	5-25
5.4.9	Creating Vertices	5-26
5.4.10	Creating Edges	5-26
5.4.11	Deleting Vertices and Edges	5-27
5.4.12	Reading a Graph from a Database into an Embedded In-Memory Analyst	5-27
5.4.13	Specifying Labels for Vertices	5-28
5.4.14	Building an In-Memory Graph	5-29
5.4.15	Dropping a Property Graph	5-30
5.4.16	Executing PGQL Queries	5-30
5.5	Managing Text Indexing for Property Graph Data	5-30
5.5.1	Configuring a Text Index for Property Graph Data	5-31
5.5.1.1	Configuring Text Indexes Using Oracle Text	5-31
5.5.2	Using Automatic Indexes for Property Graph Data	5-34
5.5.3	Using Manual Indexes for Property Graph Data	5-35
5.5.4	Executing Search Queries Over a Property Graph's Text Indexes	5-36
5.5.4.1	Executing Search Queries Over a Text Index Using Oracle Text	5-36
5.5.5	Handling Data Types	5-37
5.5.5.1	Handling Data Types on Oracle Text	5-37
5.5.6	Updating Configuration Settings on Text Indexes for Property Graph Data	5-38
5.5.7	Using Parallel Query on Text Indexes for Property Graph Data	5-38
5.5.7.1	Parallel Text Search Using Oracle Text	5-39
5.6	Access Control for Property Graph Data (Graph-Level and OLS)	5-40
5.6.1	Applying Oracle Label Security (OLS) on Property Graph Data	5-41
5.7	SQL-Based Property Graph Query and Analytics	5-46
5.7.1	Simple Property Graph Queries	5-47
5.7.2	Text Queries on Property Graphs	5-50
5.7.3	Navigation and Graph Pattern Matching	5-54
5.7.4	Navigation Options: CONNECT BY and Parallel Recursion	5-59
5.7.5	Pivot	5-62
5.7.6	SQL-Based Property Graph Analytics	5-63
5.7.6.1	Shortest Path Examples	5-64
5.7.6.2	Collaborative Filtering Overview and Examples	5-68
5.8	Creating Property Graph Views on an RDF Graph	5-73

6 Property Graph Query Language (PGQL)

6.1	Creating a Property Graph using PGQL	6-1
6.2	Pattern Matching with PGQL	6-3
6.3	Edge Patterns Have a Direction with PGQL	6-4
6.4	Vertex and Edge Labels with PGQL	6-5
6.5	Variable-Length Paths with PGQL	6-5
6.6	Aggregation and Sorting with PGQL	6-5
6.7	Supported PGQL Features and Limitations	6-6
6.8	Executing PGQL Queries Against the In-Memory Graph Server (PGX)	6-9
6.8.1	Getting Started with PGQL	6-10
6.8.2	PGQL Features Supported in the Graph Server (PGX)	6-11
6.8.2.1	Limitations on Quantifiers	6-12
6.8.2.2	Limitations on WHERE and COST Clauses in Quantified Patterns	6-12
6.8.3	Java APIs for Executing CREATE PROPERTY GRAPH Statements	6-12
6.8.4	Python APIs for Executing CREATE PROPERTY GRAPH Statements	6-13
6.8.5	Java APIs for Executing SELECT Queries	6-13
6.8.5.1	Executing SELECT Queries Against a Graph in the In-memory Graph Server (PGX)	6-14
6.8.5.2	Executing SELECT Queries Against a PGX Session	6-14
6.8.5.3	Iterating Through a Result Set	6-14
6.8.5.4	Printing a Result Set	6-17
6.8.6	Java APIs for Executing UPDATE Queries	6-18
6.8.6.1	Updatability of Graphs Through PGQL	6-18
6.8.6.2	Executing UPDATE Queries against a Graph in the in-memory Graph Server (PGX)	6-19
6.8.6.3	Executing UPDATE Queries Against a PGX Session	6-20
6.8.6.4	Altering the Underlying Schema of a Graph	6-20
6.8.7	PGQL Queries with Partitioned IDs	6-20
6.8.8	Security Tools for Executing PGQL Queries	6-22
6.8.8.1	Using Bind Variables	6-22
6.8.8.2	Using Identifiers in a Safe Manner	6-24
6.8.9	Best Practices for Tuning PGQL Queries	6-24
6.8.9.1	Memory Allocation	6-25
6.8.9.2	Parallelism	6-25
6.8.9.3	Query Plan Explaining	6-25
6.9	Executing PGQL Queries Directly Against Oracle Database	6-26
6.9.1	Executing PGQL Queries Against Property Graph Schema Tables	6-26
6.9.1.1	PGQL Features Supported in Property Graph Schema	6-27
6.9.1.2	Creating Property Graphs through CREATE PROPERTY GRAPH Statements	6-30

6.9.1.3	Dropping Property Graphs through DROP PROPERTY GRAPH Statements	6-36
6.9.1.4	Using the oracle.pg.rdbms.pgql Java Package to Execute PGQL Queries	6-38
6.9.1.5	Using the Python Client to Execute PGQL Queries	6-104
6.9.1.6	Performance Considerations for PGQL Queries	6-110
6.9.2	Executing PGQL Queries Against Property Graph Views	6-111
6.9.2.1	PGQL Features Supported in Property Graph Views	6-112
6.9.2.2	Performance Considerations for PGQL Queries	6-114
6.9.2.3	Creating a Property Graph View	6-116
6.9.2.4	Executing PGQL SELECT Queries	6-119
6.9.2.5	Dropping A Property Graph View	6-121

7 Graph Visualization Application

7.1	About the Graph Visualization Application	7-1
7.2	How does the Graph Visualization Application Work	7-1
7.3	Using the Graph Visualization Application	7-2
7.3.1	Graph Visualization Modes	7-3
7.3.2	Graph Visualization Settings	7-3
7.3.3	Using the Geographical Layout	7-6
7.3.4	Using Live Search	7-8
7.3.5	Using URL Parameters to Control the Graph Visualization Application	7-9
7.4	REST Endpoints for the Graph Visualization Application	7-9
7.4.1	Login	7-10
7.4.2	List Graphs	7-11
7.4.3	Run a PGQL Query	7-11
7.4.4	Get User	7-13
7.4.5	Asynchronous REST Endpoints	7-14
7.4.5.1	Run a PGQL Query Asynchronously	7-14
7.4.5.2	Check a Query Completion	7-14
7.4.5.3	Cancel a Query Execution	7-15
7.4.5.4	Retrieve a Query Result	7-15
7.5	Kerberos Enabled Authentication for the Graph Visualization Application	7-17
7.5.1	Prerequisite Requirements for Kerberos Authentication	7-17
7.5.2	Preparing the Graph Visualization Application for Kerberos Authentication	7-18

8 Using the Machine Learning Library (PgxML) for Graphs

8.1	Using the DeepWalk Algorithm	8-1
8.1.1	Loading a Graph	8-2
8.1.2	Building a Minimal DeepWalk Model	8-3
8.1.3	Building a Customized DeepWalk Model	8-4

8.1.4	Training a DeepWalk Model	8-5
8.1.5	Getting the Loss Value For a DeepWalk Model	8-6
8.1.6	Computing Similar Vertices for a Given Vertex	8-6
8.1.7	Computing Similar Vertices for a Vertex Batch	8-7
8.1.8	Getting All Trained Vertex Vectors	8-9
8.1.9	Storing a Trained DeepWalk Model	8-9
8.1.9.1	Storing a Trained Model in Another Database	8-10
8.1.10	Loading a Pre-Trained DeepWalk Model	8-12
8.1.10.1	Loading a Pre-Trained Model From Another Database	8-13
8.1.11	Destroying a DeepWalk Model	8-14
8.2	Using the Supervised GraphWise Algorithm	8-15
8.2.1	Loading a Graph	8-16
8.2.2	Building a Minimal GraphWise Model	8-17
8.2.3	Advanced Hyperparameter Customization	8-19
8.2.4	Setting a Custom Loss Function and Batch Generator (for Anomaly Detection)	8-21
8.2.5	Training a Supervised GraphWise Model	8-22
8.2.6	Getting the Loss Value For a Supervised GraphWise Model	8-23
8.2.7	Inferring the Vertex Labels for a Supervised GraphWise Model	8-23
8.2.8	Evaluating the Supervised GraphWise Model Performance	8-25
8.2.9	Inferring Embeddings for a Supervised GraphWise Model	8-26
8.2.9.1	Inferring Embeddings for a Model in Another Database	8-27
8.2.10	Storing a Trained Supervised GraphWise Model	8-28
8.2.11	Loading a Pre-Trained Supervised GraphWise Model	8-29
8.2.12	Destroying a Supervised GraphWise Model	8-30
8.2.13	Explaining a Prediction of a Supervised GraphWise Model	8-31
8.3	Using the Unsupervised GraphWise Algorithm	8-34
8.3.1	Loading a Graph	8-35
8.3.2	Building a Minimal Unsupervised GraphWise Model	8-36
8.3.3	Advanced Hyperparameter Customization	8-37
8.3.4	Training a Unsupervised GraphWise Model	8-38
8.3.5	Getting the Loss Value for a Unsupervised GraphWise Model	8-39
8.3.6	Inferring Embeddings for a Unsupervised GraphWise Model	8-39
8.3.7	Storing a Unsupervised GraphWise Model	8-41
8.3.8	Loading a Pre-Trained Unsupervised GraphWise Model	8-42
8.3.9	Destroying a Unsupervised GraphWise Model	8-43
8.3.10	Explaining a Prediction for an Unsupervised GraphWise Model	8-43
8.4	Using the Pg2vec Algorithm	8-46
8.4.1	Loading a Graph	8-47
8.4.2	Building a Minimal Pg2vec Model	8-49
8.4.3	Building a Customized Pg2vec Model	8-50
8.4.4	Training a Pg2vec Model	8-51

8.4.5	Getting the Loss Value For a Pg2vec Model	8-52
8.4.6	Computing Similar Graphlets for a Given Graphlet	8-52
8.4.7	Computing Similar for a Graphlet Batch	8-54
8.4.8	Inferring a Graphlet Vector	8-55
8.4.9	Inferring Vectors for a Graphlet Batch	8-56
8.4.10	Storing a Trained Pg2vec Model	8-57
8.4.11	Loading a Pre-Trained Pg2vec Model	8-58
8.4.12	Destroying a Pg2vec Model	8-59

9 OPG_APIS Package Subprograms

9.1	OPG_APIS.ANALYZE_PG	9-2
9.2	OPG_APIS.CF	9-4
9.3	OPG_APIS.CF_CLEANUP	9-7
9.4	OPG_APIS.CF_PREP	9-9
9.5	OPG_APIS.CLEAR_PG	9-10
9.6	OPG_APIS.CLEAR_PG_INDICES	9-11
9.7	OPG_APIS.CLONE_GRAPH	9-11
9.8	OPG_APIS.COUNT_TRIANGLE	9-12
9.9	OPG_APIS.COUNT_TRIANGLE_CLEANUP	9-13
9.10	OPG_APIS.COUNT_TRIANGLE_PREP	9-14
9.11	OPG_APIS.COUNT_TRIANGLE_RENUM	9-16
9.12	OPG_APIS.CREATE_EDGES_TEXT_IDX	9-17
9.13	OPG_APIS.CREATE_PG	9-18
9.14	OPG_APIS.CREATE_PG_SNAPSHOT_TAB	9-19
9.15	OPG_APIS.CREATE_PG_TEXTIDX_TAB	9-21
9.16	OPG_APIS.CREATE_STAT_TABLE	9-22
9.17	OPG_APIS.CREATE_SUB_GRAPH	9-23
9.18	OPG_APIS.CREATE_VERTICES_TEXT_IDX	9-24
9.19	OPG_APIS.DROP_EDGES_TEXT_IDX	9-26
9.20	OPG_APIS.DROP_PG	9-26
9.21	OPG_APIS.DROP_PG_VIEW	9-27
9.22	OPG_APIS.DROP_VERTICES_TEXT_IDX	9-27
9.23	OPG_APIS.ESTIMATE_TRIANGLE_RENUM	9-28
9.24	OPG_APIS.EXP_EDGE_TAB_STATS	9-30
9.25	OPG_APIS.EXP_VERTEX_TAB_STATS	9-31
9.26	OPG_APIS.FIND_CC_MAPPING_BASED	9-32
9.27	OPG_APIS.FIND_CLUSTERS_CLEANUP	9-33
9.28	OPG_APIS.FIND_CLUSTERS_PREP	9-34
9.29	OPG_APIS.FIND_SP	9-36
9.30	OPG_APIS.FIND_SP_CLEANUP	9-37

9.31	OPG_APIS.FIND_SP_PREP	9-38
9.32	OPG_APIS.GET_BUILD_ID	9-39
9.33	OPG_APIS.GET_GEOMETRY_FROM_V_COL	9-39
9.34	OPG_APIS.GET_GEOMETRY_FROM_V_T_COLS	9-41
9.35	OPG_APIS.GET_LATLONG_FROM_V_COL	9-42
9.36	OPG_APIS.GET_LATLONG_FROM_V_T_COLS	9-43
9.37	OPG_APIS.GET_LONG_LAT_GEOMETRY	9-44
9.38	OPG_APIS.GET_LATLONG_FROM_V_COL	9-45
9.39	OPG_APIS.GET_LONGLAT_FROM_V_T_COLS	9-46
9.40	OPG_APIS.GET_SCN	9-47
9.41	OPG_APIS.GET_VERSION	9-47
9.42	OPG_APIS.GET_WKTGEOMETRY_FROM_V_COL	9-48
9.43	OPG_APIS.GET_WKTGEOMETRY_FROM_V_T_COLS	9-49
9.44	OPG_APIS.GRANT_ACCESS	9-50
9.45	OPG_APIS.IMP_EDGE_TAB_STATS	9-51
9.46	OPG_APIS.IMP_VERTEX_TAB_STATS	9-52
9.47	OPG_APIS.PR	9-54
9.48	OPG_APIS.PR_CLEANUP	9-56
9.49	OPG_APIS.PR_PREP	9-57
9.50	OPG_APIS.PREPARE_TEXT_INDEX	9-58
9.51	OPG_APIS.RENAME_PG	9-58
9.52	OPG_APIS.SPARSIFY_GRAPH	9-59
9.53	OPG_APIS.SPARSIFY_GRAPH_CLEANUP	9-61
9.54	OPG_APIS.SPARSIFY_GRAPH_PREP	9-62

10 OPG_GRAPHOP Package Subprograms

10.1	OPG_GRAPHOP.POPULATE_SKELETON_TAB	10-1
------	-----------------------------------	------

Part II In-Memory Graph Server (PGX) Advanced User Guide

11 Configuring the In-Memory Graph Server (PGX)

11.1	Configuration Parameters for the Graph Server (PGX) Engine	11-1
11.1.1	Configuration of the Graph Server (PGX) Run-Time Parameters	11-11
11.1.2	Specifying the Configuration File to the In-Memory Graph Server (PGX)	11-14
11.1.3	Memory Consumption by the Graph Server (PGX)	11-15
11.1.3.1	Memory Management	11-15
11.2	Configuration Parameters for Connecting to the Graph Server (PGX)	11-17
11.3	Configuration Parameters for the Graph Client	11-17

12 Graphs Management

12.1	Loading a Graph Into the Graph Server (PGX)	12-1
12.1.1	API for Loading Graphs into Memory	12-1
12.1.2	Graph Configuration Options	12-2
12.1.3	Preloading a Graph	12-10
12.1.4	Data Loading Security Best Practices	12-10
12.1.5	Data Format Support Matrix	12-11
12.1.6	Immutability of Loaded Graphs	12-11
12.2	Publishing a Graph	12-12
12.3	Publishing a Preloaded Graph	12-17
12.4	Deleting a Graph	12-18

13 Namespaces and Sharing

13.1	Defining Graph Names	13-1
13.2	Retrieving Graphs by Name	13-1
13.3	Checking Used Names	13-2
13.4	Property Name Resolution and Graph Mutations	13-2

14 PGX Programming Guides

14.1	Design of the Graph Server (PGX) API	14-3
14.2	Data Types and Collections in the Graph Server (PGX)	14-4
14.2.1	Using Collections and Maps	14-7
14.2.1.1	Collection Data Types	14-7
14.2.1.2	Map Data Types	14-12
14.2.2	Using Datetime Data Types	14-17
14.2.2.1	Loading Datetime Data	14-18
14.2.2.2	Specifying Custom Datetime Formats	14-20
14.2.2.3	APIs for Accessing Datetime Data	14-21
14.2.2.4	Querying Datetime Data Using PGQL	14-22
14.2.2.5	Accessing Datetimes from PGQL Result Sets	14-24
14.3	Handling Asynchronous Requests in Graph Server (PGX)	14-26
14.3.1	Blocking Operation	14-26
14.3.2	Chaining Operation	14-27
14.3.3	Cancelling Operation	14-28
14.3.4	Handling Concurrent Asynchronous Operations	14-28
14.4	Graph Client Sessions	14-29
14.5	Graph Mutation and Subgraphs	14-31
14.5.1	Altering Graphs	14-31
14.5.1.1	Loading Or Removing Additional Vertex or Edge Providers	14-32

14.5.2	Simplifying and Copying Graphs	14-40
14.5.3	Transposing Graphs	14-42
14.5.4	Undirecting Graphs	14-43
14.5.5	Advanced Multi-Edge Handling	14-44
14.5.5.1	Picking	14-44
14.5.5.2	Merging	14-45
14.5.5.3	StrategyBuilder in General	14-46
14.5.6	Creating a Subgraph	14-47
14.5.7	Creating a Bipartite Subgraph	14-47
14.5.8	Creating a Sparsified Subgraph	14-48
14.6	Graph Builder and Graph Change Set	14-49
14.6.1	Building Graphs Using GraphBuilder Interface	14-49
14.6.1.1	Creating a Simple Graph	14-50
14.6.1.2	Adding a Vertex Property	14-51
14.6.1.3	Using Strings as Vertex Identifiers	14-53
14.6.1.4	Referencing a Vertex for Creating Edges	14-54
14.6.1.5	Adding an Edge Property and a Label	14-56
14.6.1.6	Using Graph Builder with Implicit IDs	14-57
14.6.2	Modifying Loaded Graphs Using ChangeSet	14-59
14.6.2.1	Modifying Vertices	14-59
14.6.2.2	Adding Edges	14-60
14.6.2.3	GraphChangeSet with Partitioned IDs	14-61
14.6.2.4	Error Handling when Using a ChangeSet	14-62
14.7	Managing Transient Data	14-63
14.7.1	Managing Transient Properties	14-64
14.7.2	Managing Collections and Scalars	14-65
14.8	Graph Versioning	14-67
14.8.1	Configuring the Snapshots Source	14-68
14.8.2	Creating a Snapshot via Refreshing	14-68
14.8.3	Creating a Snapshot via ChangeSet	14-71
14.8.4	Checking Out the Latest Snapshots of a Graph	14-73
14.8.5	Checking Out Different Snapshots of a Graph	14-74
14.8.6	Directly Loading a Specific Snapshot of a Graph	14-75
14.9	Labels and Properties	14-77
14.9.1	Setting and Getting Property Values	14-77
14.9.2	Getting Label Values	14-79
14.10	Filter Expressions	14-79
14.10.1	Syntax	14-80
14.10.2	Type System	14-84
14.10.3	Path Finding Filters	14-85
14.10.4	Subgraph Filters	14-85

14.10.5	Operations on Filter Expressions	14-86
14.10.5.1	Defining Filter Expressions	14-86
14.10.5.2	Defining Result Set Filters	14-87
14.10.5.3	Creating a Subgraph from PGQL Result Set	14-88
14.10.5.4	Defining Collection Filters	14-90
14.10.5.5	Creating a Subgraph from Collection Filters	14-91
14.10.5.6	Combining Filter Expressions	14-92
14.10.5.7	Creating a Subgraph Using Filter Expressions with Partitioned IDs	14-93
14.11	Advanced Task Scheduling Using Execution Environments	14-95
14.11.1	Enterprise Scheduler Configuration Guide	14-95
14.11.2	Enabling Enterprise Scheduler Features	14-97
14.11.3	Retrieving and Inspecting the Execution Environment	14-97
14.11.4	Modifying and Submitting Tasks Under an Updated Environment	14-99
14.11.5	Using Lambda Syntax	14-100
14.12	Admin API	14-101
14.12.1	Get a Server Instance	14-102
14.12.2	Get Inspection Data	14-102
14.12.3	Get Active Sessions	14-103
14.12.4	Get Cached Graphs	14-105
14.12.5	Get Published Graphs	14-106
14.12.6	Get Currently Loading Graphs	14-106
14.12.7	Get Tasks	14-107
14.12.8	Get Available Memories	14-107
14.13	PgxFrames Tabular Data-Structure	14-107
14.13.1	Loading a PgxFrame from a Database	14-108
14.13.2	Loading a PgxFrame from Client-Side Data	14-111
14.13.3	Printing the Content of a PgxFrame	14-115
14.13.4	Destroying a PgxFrame	14-116
14.13.5	Storing a PgxFrame to a Database	14-116
14.13.6	Loading and Storing Vector Properties	14-118
14.13.7	Flattening Vector Properties	14-119
14.13.8	Union of PGX Frames	14-120
14.13.9	Joining PGX Frames	14-121
14.13.10	PgxFrame Helpers	14-122
14.13.11	PgxFrame-PgqlResultSet Conversions	14-125
14.13.12	Creating a Graph from Multiple PgxFrame Objects	14-127

15 Working with Files Using the Graph Server (PGX)

15.1	Loading Graph Data from Files	15-1
15.1.1	Graph Configuration for Loading from File	15-2

15.1.2	Specifying the File Path	15-7
15.1.3	Supported File Access Protocols	15-7
15.1.4	Plain Text Formats	15-8
15.1.4.1	Comma-Separated Values (CSV)	15-10
15.1.4.2	Adjacency List (ADJ_LIST)	15-13
15.1.4.3	Edge List (EDGE_LIST)	15-13
15.1.4.4	Two Tables (TWO_TABLES)	15-15
15.1.5	XML File Formats	15-16
15.1.6	Binary File Formats	15-17
15.2	Loading Graph Data in Parallel from Multiple Files	15-22
15.3	Exporting Graphs Into a File	15-24
15.3.1	Exporting a Graph to Disk	15-26
15.4	Exporting a Graph into Multiple Files	15-28

16 Log Management in the Graph Server (PGX)

16.1	Configuring Log4j Logging	16-1
------	---------------------------	------

Part III Supplementary Information for Property Graph Support

A Mapping Graph Server Roles to Default Privileges

B Disabling Transport Layer Security (TLS) in Graph Server

Index

List of Figures

1-1	Simple Property Graph Example	1-3
1-2	Three-Tier Property Graph Architecture	1-5
1-3	Two-Tier Property Graph Architecture	1-6
1-4	Graph Visualization Login	1-20
1-5	PGQL on Graph Server (PGX)	1-23
1-6	PGQL on Database	1-24
1-7	Enabling Accessibility Mode in the Graph Visualization Application	1-46
3-1	Financial Transactions Graph	3-6
4-1	Graph Server (PGX) Design	4-3
4-2	Session and Transient Properties	4-4
4-3	Remote Server Mode	4-5
4-4	PGX as a Library	4-5
4-5	Edges Matching <code>src.prop == 10</code>	4-56
4-6	Graph Created by the Simple Filter	4-56
4-7	Edges Matching the <code>outDegree</code> Filter	4-57
4-8	Graph Created by the <code>outDegree</code> Filter	4-58
5-1	Phones Graph for Collaborative Filtering	5-68
6-1	PGQL on Property Graph Schema Tables in Oracle Database (RDBMS)	6-27
6-2	PGQL on Property Graph Views in Oracle Database	6-111
7-1	Query Visualization	7-2
7-2	Graph Visualization Settings Window	7-4
7-3	Highlights Options for Vertices	7-5
7-4	Geographical Layout	7-6
7-5	Setting Geographical Layout	7-7
7-6	Selecting the Coordinates for the Geographical layout	7-8
8-1	Pg2vec - Visualization of Two Similar Graphlets	8-54
14-1	Picking Strategy	14-45
14-2	Merging Strategy	14-46

List of Tables

1-1	Graph Size Estimator	1-5
1-2	Overview of Tasks to Get Started with Property Graphs	1-7
1-3	Components in the Oracle Graph Server and Client Deployment	1-8
3-1	Metadata Tables for PG Views	3-1
4-1	Privileges and Roles in Oracle Database	4-7
4-2	Advanced Access Configuration Options	4-17
4-3	API for Checking Graph Permissions	4-18
4-4	Allowed Permissions	4-22
4-5	Overview of Built-In Algorithms	4-46
4-6	Configuration Parameters for the In-Memory Graph Server (PGX)	4-66
4-7	Fields for Each UDF	4-84
6-1	CREATE PROPERTY GRAPH Statement Support	6-3
6-2	Supported PGQL Functionalities and Limitations	6-6
6-3	Type Casting Support in PGQL (From and To Types)	6-29
6-4	PGQL Translation and Execution Options	6-69
6-5	PGQL Statement Modification Options	6-100
7-1	Available URL Parameters	7-9
7-2	Parameters	7-10
7-3	Query Parameters	7-11
7-4	Location of WEB-INF/web.xml file	7-18
11-1	Configuration Parameters for the Graph Server (PGX) Engine	11-2
11-2	Graph Server (PGX) Run-Time Parameters	11-11
11-3	Configuration Parameters for the Graph Client	11-17
12-1	Graph Config JSON Fields	12-2
12-2	Provider Configuration JSON file Options	12-5
12-3	Property Configuration	12-6
12-4	Loading Configuration	12-8
12-5	Error Handling Configuration	12-9
12-6	Data Format Support Matrix	12-11
14-1	PGX API Interface	14-1
14-2	Overview of Data types	14-5
14-3	Overview of Datetime Data Types in PGX	14-18
14-4	Default Property Values	14-52
14-5	Default Temporal Formats	14-81
14-6	Session Information Options	14-104

14-7	Graph Information	14-105
14-8	Mapping between In-Place and Out-Place Operations	14-107
15-1	Loading from File - Graph Configuration Options	15-2
15-2	CSV Specific Options	15-5
15-3	Type Encoding	15-17
15-4	File Layout	15-17
15-5	Integer Vertex Keys	15-19
15-6	Long Vertex Keys	15-19
15-7	String Vertex Keys	15-19
15-8	String Key Element Layout	15-19
15-9	Primitive Type Layout	15-19
15-10	Vector Property Layout	15-20
15-11	String Type Layout	15-20
15-12	String Dictionary Layout	15-20
15-13	String Dictionary Element Layout	15-20
15-14	Vertex Labels Layout	15-21
15-15	Shared Pools Layout	15-21
15-16	Type == Enum	15-21
15-17	Type == Prefix	15-22
15-18	String Table for Shared Pools	15-22
15-19	Property Names Layout	15-22
15-20	Files CompressionScheme	15-25
15-21	Graph Configuration when Exporting Graph into Multiple Files	15-25
A-1	Mapping Graph Server Roles to Default Privileges	A-1

Preface

This document provides conceptual and usage information about Oracle Database support for working with property graph data.

- [Audience](#)
- [Documentation Accessibility](#)
- [Related Documents](#)
- [Conventions](#)

Audience

This document is intended for database and application developers in an Oracle Database environment.

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.

Related Documents

For more information, see the following documents:

- *Oracle Spatial Developer's Guide*
- *Oracle Database Graph Developer's Guide for RDF Graph*
- *Oracle Spatial GeoRaster Developer's Guide*
- *Oracle Spatial Topology and Network Data Model Developer's Guide*
- *Oracle Big Data Spatial and Graph User's Guide and Reference*

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.

Changes in This Release for This Guide

The following changes apply to property graph support that is shipped with Oracle Graph Server and Client.

Oracle Graph Server and Client is required for using the property graph feature of Oracle Database (see [Oracle Graph Server and Client Installation](#)), and is released four times a year.

New Features

Significant New Features in Oracle Graph Server and Client 21.4

- Enhanced support for `GraphAlterationBuilder` API to add empty vertex or edge providers that are not linked to any datasource.
See [Adding Additional Empty Vertex or Edge Providers](#) for more details.
- Added Oracle functions and user defined functions (UDFs) support in PGQL queries for property graph views.
See [PGQL Features Supported in Property Graph Views](#) for more information.
- Enhanced user friendly error reporting based on graph metadata by the PGQL parser when validating PGQL queries executed against the graph server (PGX).
- Added `GNNExplainer` support to leverage Graph ML explainability for Unsupervised GraphWise models in PgxML Library.
See [Explaining a Prediction for an Unsupervised GraphWise Model](#) for more information.
- Added support for custom loss functions and batch generators for anomaly detection with Supervised GraphWise models.
See [Setting a Custom Loss Function and Batch Generator \(for Anomaly Detection\)](#) for more information.

Migrating Property Graph Applications to Oracle Database 21c

From Release 21c onwards, Oracle Graph Server and Client must be installed separately. It is recommended to remove the older property graph libraries from `$ORACLE_HOME`. See [Uninstalling Previous Versions of Property Graph Libraries](#) for more details.

Deprecated Features

- **PL/SQL API `OPG_APIS.GET_SCN` Function**
The PL/SQL API `OPG_APIS.GET_SCN` function is deprecated. Instead, to retrieve the current SCN (system change number), use the `DBMS_FLASHBACK.GET_SYSTEM_CHANGE_NUMBER` function:

```
SELECT dbms_flashback.get_system_change_number FROM DUAL;
```

- **OraclePgqlResultSet**

The `oracle.pg.rdbms.OraclePgqlResultSet` interface was deprecated in 19c and will be removed in a future release. Instead, use the standardized interface `oracle.pgql.lang.ResultSet` to retrieve values from a PGQL result set.

Desupported Features

- The following Java API methods, objects and fields in `oracle.pgx.api` are no longer supported:

Desupported Methods:

- **PgxCollection methods:**
 - * `addAllAsync(Collection<E> source)`
 - * `removeAllAsync(Collection<E> source)`
 - * `addAll(ID...ids)`
 - * `removeAll(ID...ids)`
- **PgqlResultSet methods:**
 - * `getResults()`: instead, use `PgqlResultSet` to directly iterate the result set
 - * `destroy()`
- **User-defined pattern matching semantic methods:**
 - * `PgxGraph#queryPgql(String, PatternMatchingSemantic)`: instead, use `PgxGraph#queryPgql(String)`
 - * `PgxSession.setPatternMatchingSemantic(..)`
- **GraphMetaData constructors and related methods:**
 - * `GraphMetaData()`
 - * `GraphMetaData(GraphMetaData other, java.net.URI baseUri)`
 - * `GraphMetaData(IdType vertexIdType)`
 - * `GraphMetaData.setVertexIdType()`
 - * `GraphMetaData.setEdgeIdType()`
- `PgxSession#getAvailableSnapshots(GraphConfig)`: instead, use `PgxSession#getAvailableSnapshots(PgxGraph)`
- **All `Analyst#filteredBfs` and `Analyst#filteredDfs` methods that accepts `filter` parameter: instead, use the `navigator` parameter**

Desupported Objects

- `PgqlResult` (a result of `resultSet.getResults().iterator().next()`): instead, use `PgxResult` as returned from `resultSet.iterator().next()`

Desupported Fields

- `pattern_matching_semantic` configuration field
- The Java API method `AbstractGraphConfigBuilder#setNodeIdType` in `oracle.pgx.config` is desupported. Instead, use `AbstractGraphConfigBuilder#setVertexIdType()`.
- The following PyPGX classes are desupported in `pypgx` package. Instead, use `pypgx.api.filters` subpackage to access these classes:

- `EdgeFilter`
- `GraphFilter`
- `VertexFilter`
- The following PyPGX classes are desupported in `pypgx.api` package. Instead, use `pypgx.api.frames` subpackage to access these classes:
 - `PgxCsvFrameReader`
 - `PgxCsvFrameStorer`
 - `PgxDbFrameReader`
 - `PgxDbFrameStorer`
 - `PgxFrame`
 - `PgxFrameBuilder`
 - `PgxFrameColumn`
 - `PgxGenericFrameReader`
 - `PgxGenericFrameStorer`
 - `PgxPgbFrameReader`
 - `PgxPgbFrameStorer`
- The following Python API packages are no longer supported:
 - `common`: This internal package is desupported. Few of the classes from this package are moved to the public package `pypgx.api`.
 - `utils`: This internal package is renamed to `_utils`.
- Graph property text search based on Apache Solr/Lucene is desupported. Instead, use Oracle Text or PGQL query expressions.
- The PGX property type `DATE` is desupported. Instead, use `LOCAL_DATE` or `TIMESTAMP`.
- Property Graph support for data stored in Oracle NoSQL Database is desupported.
- The two-table format is desupported.
- The Gremlin Groovy shell is desupported.
- Apache Tinkerpop API support for Oracle Database is desupported.
- Loading data from flat file formats into the property graph schema is desupported.
- Apache Groovy-based shell was deprecated in 19c and is now desupported.
- Apache HBASE and Apache HDFS on Cloudera CDH5 is desupported.

Part I

Getting Started with Oracle Property Graphs

Part I provides the fundamental information to get you started on the property graph feature of Oracle Database.

This part covers the following:

- Overview of the property graph features
- Installing and configuring Oracle Graph Server and Client
- Graph data modeling using the in-memory graph server (PGX)
- Querying graph data using PGQL
- Visualizing property graphs using the Graph Visualization Application
- Applying machine learning algorithms to analyze relationships in graph data

Part I contains the following chapters:

- [Property Graph Support Overview](#)
- [Quick Starts for Using Oracle Property Graph](#)
This chapter contains quick start tutorials and other resources to help you get started on working with Oracle property graphs.
- [Property Graph Views on Oracle Database Tables](#)
You can create property graph views over data stored in Oracle Database. You can perform various graph analytics operations using PGQL on these views.
- [Using the In-Memory Graph Server \(PGX\)](#)
The in-memory Graph server of Oracle Graph supports a set of analytical functions.
- [Using the Property Graph Schema](#)
This chapter provides conceptual and usage information about creating, storing, and working with property graph data in an Oracle Database environment.
- [Property Graph Query Language \(PGQL\)](#)
PGQL is a SQL-like query language for property graph data structures that consist of *vertices* that are connected to other vertices by *edges*, each of which can have key-value pairs (properties) associated with them.
- [Graph Visualization Application](#)
The Graph Visualization application enables interactive exploration and visualization of property graphs. It can also visualize graphs stored in the database.
- [Using the Machine Learning Library \(PgxML\) for Graphs](#)
The in-memory graph server (PGX) provides a machine learning library `oracle.pgx.api.ml.lib`, which supports graph-empowered machine learning algorithms.
- [OPG_APIS Package Subprograms](#)
The OPG_APIS package contains subprograms (functions and procedures) for working with property graphs in an Oracle database.
- [OPG_GRAPHOP Package Subprograms](#)
The OPG_GRAPHOP package contains subprograms for various operations on property graphs in an Oracle database.

1

Property Graph Support Overview

This chapter provides an overview of Oracle Graph support for property graph features.

- [Introduction to Property Graphs](#)
Property graphs give you a different way of looking at your data.
- [Before You Begin with Oracle Property Graphs](#)
Before you begin to create an Oracle Property Graph, you may need to adhere to one or more of the prerequisites explained in this section.
- [Oracle Graph Server and Client Installation](#)
This section explains the various operations that you must perform to install, uninstall or upgrade Oracle Graph Server (PGX). It also includes the steps to install Oracle Graph Client.
- [Setting Up Transport Layer Security](#)
The graph server (PGX), by default, allows only encrypted connections using Transport Layer Security (TLS). TLS requires the server to present a server certificate to the client and the client must be configured to trust the issuer of that certificate.
- [Adding Permissions to Publish the Graph](#)
There are two ways by which you can view any graph in your graph server (PGX) session in the graph visualization application.
- [Security Best Practices with Graph Data](#)
Several security-related best practices apply when working with graph data.
- [Interactive Graph Shell CLIs](#)
Both the Oracle Graph server and client packages contain interactive command-line applications for interacting with the Java APIs and the Python APIs of the product, locally or on remote computers.
- [Developing Applications Using Graph Server Functionality as a Library](#)
The graph functions available with the graph server (PGX) can be used as a library in your application.
- [Storing Graphs in Oracle Database and Loading Graphs into Memory](#)
You can work with graphs in **two-tier mode** (graph client connects directly to Oracle Database), or **three-tier mode** (graph client connects to the graph server (PGX) on the middle-tier, which then connects to Oracle Database).
- [Using Oracle Graph with the Autonomous Database](#)
Oracle Graph with the Autonomous Database allows you to create property graphs from data in your Autonomous Database.
- [Migrating Property Graph Applications from Before Release 21c](#)
If you are migrating from a previous version of Oracle Spatial and Graph to Release 21c, you may need to make some changes to existing property graph-related applications.
- [Upgrading From Graph Server and Client 20.4.x to 21.x](#)
If you are upgrading from Graph Server and Client 20.4.x to 21.x version, you may need to create new roles in database and migrate authorization rules from `pgx.conf` file to the database. Also, starting from Graph Server and Client Release 21.1, TLS is enforced at the time of the RPM file installation.

- [Using the Graph Zeppelin Interpreter Client](#)
Oracle Graph provides an interpreter client implementation for Apache Zeppelin. This tutorial topic explains how to perform simple operations using the graph Zeppelin interpreter client.
- [About Oracle Graph Server and Client Accessibility](#)
This section provides information on the accessibility features for Oracle Graph Server and Client.

1.1 Introduction to Property Graphs

Property graphs give you a different way of looking at your data.

You can model your data as a graph by making data entities **vertices** in the graph, and relationships between them as **edges** in the graph. For example, in a bank customer accounts can be vertices, and cash transfer relationships between them can be edges.

When you view your data as a graph, you can analyze your data based on the connections and relationships between them. You can run graph analytics algorithms like PageRank to measure the relative importance of data entities based on the relationships between them, for example, links between webpages.

- [What Are Property Graphs?](#)
- [About the Property Graph Feature of Oracle Database](#)
The Property Graph feature delivers advanced graph query and analytics capabilities in Oracle Database.

1.1.1 What Are Property Graphs?

A property graph consists of a set of objects or **vertices**, and a set of arrows or **edges** connecting the objects. Vertices and edges can have multiple properties, which are represented as key-value pairs.

Each vertex has a unique identifier and can have:


- A set of outgoing edges
- A set of incoming edges
- A collection of properties

Each edge has a unique identifier and can have:

- An outgoing vertex
- An incoming vertex
- A text label that describes the relationship between the two vertices
- A collection of properties

For vertices and edges, each property is identified with a unique name.

The following figure illustrates a very simple property graph with two vertices and one edge. The two vertices have identifiers 1 and 2. Both vertices have properties `name` and `age`. The edge is from the outgoing vertex 1 to the incoming vertex 2. The edge has a text label `knows` and a property `type` identifying the type of relationship between vertices 1 and 2.

Figure 1-1 Simple Property Graph Example

A property graph can have self-edges (that is, an edge whose source and destination vertex are the same), as well as multiple edges between the same source and destination vertices.

A property graph can also have different types of vertices and edges in the same graph. For example a graph can have a set of vertices with label `Person` and a set of vertices with label `Place`, with different properties relevant to these two sets of vertices.

The property graph data model is similar to the W3C standards-based Resource Description Framework (RDF) graph data model; however, the property graph data model is simpler and less precise than RDF.

The property graph data model features and analytic APIs make property graphs a good candidate for use cases such as these:

- Identifying influencers in a social network
- Predicting trends and customer behavior
- Discovering relationships based on pattern matching
- Identifying clusters to customize campaigns

 **Note:**

The property graph data model that Oracle supports at the database side does not allow labels for vertices. However, you can treat the value of a designated vertex property as one or more labels.

Related Topics

- [Specifying Labels for Vertices](#)

1.1.2 About the Property Graph Feature of Oracle Database

The Property Graph feature delivers advanced graph query and analytics capabilities in Oracle Database.

This feature supports graph operations, indexing, queries, search, and in-memory analytics.

Graphs manage networks of linked data as vertices, edges, and properties of the vertices and edges. Graphs are commonly used to model, store, and analyze relationships found in social networks, cybersecurity, utilities and telecommunications, life sciences and clinical data, and knowledge networks.

Typical graph analyses encompass graph traversal, recommendations, finding communities and influencers, and pattern matching. Industries including telecommunications, life sciences and healthcare, security, media, and publishing can benefit from graphs.

The property graph features of Oracle Special and Graph support those use cases with the following capabilities:

- A scalable graph database
- Developer-based APIs based upon PGQL and Java graph APIs
- Text search and query through integration with Oracle Text
- A parallel, in-memory graph server (PGX) for running graph queries and graph analytics
See [In-Memory Graph Server \(PGX\)](#) for more information.
- A fast, scalable suite of social network analysis functions that include ranking, centrality, recommender, community detection, and path finding
- Parallel bulk load and export of property graph data in Oracle-defined flat files format
- A powerful Graph Visualization (GraphViz) application
- Notebook support through integration with Apache Zeppelin and Jupyter
- [In-Memory Graph Server \(PGX\)](#)
- [Options for Property Graph Architecture](#)

1.1.2.1 In-Memory Graph Server (PGX)

The in-memory graph server layer enables you to analyze property graphs using parallel in-memory execution. It provides over 60 analytic functions. Examples of the categories and specific functions include:

- Centrality - Degree Centrality, Eigenvector Centrality, PageRank, Betweenness Centrality, Closedness Centrality
- Component and Community - Strongly Connected Components (Tarjan's and Kosaraju's). Weakly Connected Components
- Twitter's Who-To-Follow, Label Propagation.
- Path Finding - Single source all destination (Bellman-Ford), Dijkstra's shortest path, Hop Distance (Breadth-first search)
- Community Evaluation - Coefficient (Triangle Counting), Conductance, Modularity, Adamic-Adar counter.

See [Using the In-Memory Graph Server \(PGX\)](#) for more information on the in-memory graph server (PGX).

1.1.2.2 Options for Property Graph Architecture

You have two architecture options when using the property graph feature of Oracle Database:

- [Run Graph Query and Analytics in the In-Memory Graph Server \(PGX\) \(3-Tier\)](#)
- [Load the Graph into Oracle Database \(2-Tier\)](#)


Both options let you use the Property Graph Query Language (PGQL).

Run Graph Query and Analytics in the In-Memory Graph Server (PGX) (3-Tier)

You can load your property graph into the in-memory graph server, which has a specialized architecture for graph computations. All query and analytics operations on this graph can be executed in-memory in the graph server. This graph can be created directly from relational tables or loaded from the property graph schema that stores the graph in the database. You can modify the graph in memory (insert, update, and delete vertices and edges, and create new properties for results of executing an algorithm). The graph server does not write the modifications back to the relational tables.

The in-memory graph server (PGX) typically in a server separate from the database, and can run standalone, or in a container like Oracle WebLogic Server or Apache Tomcat. This approach (load your property graph into the in-memory graph server) uses a three-tier architecture, as shown in the following figure.

Figure 1-2 Three-Tier Property Graph Architecture



Property Graph Sizing Recommendations

You can compute the memory required by the in-memory graph server (PGX) by using this calculator, [Graph Size Estimator](#).

For example, the following table shows the memory estimated by the calculator for the given input:

Table 1-1 Graph Size Estimator

Number of vertices	Number of Edges	Properties per Vertex	Properties per Edge	Estimated graph size
10M	100M	<ul style="list-style-type: none"> 4 - Integer Type 1 - String Type (15 characters) 	<ul style="list-style-type: none"> 4 - Integer Type 1 - String Type (15 characters) 	15 GB

Table 1-1 (Cont.) Graph Size Estimator

Number of vertices	Number of Edges	Properties per Vertex	Properties per Edge	Estimated graph size
100M	1B	<ul style="list-style-type: none"> 4 - Integer Type 1 - String Type (15 characters) 	<ul style="list-style-type: none"> 4 - Integer Type 1 - String Type (15 characters) 	140 GB

 **Note:**

- Reading a graph into memory can take up to twice the amount of memory needed to represent it in memory. So when you calculate the memory required for running PGX it is recommended that you double the amount of memory of the estimated graph size.
- **CPU Processors:** The recommended number of CPU processors for a graph with 10M vertices and 100M edges is 2-4 processors, and up to 16 processors for more compute-intensive workloads. Increasing CPU processors will improve performance.


Load the Graph into Oracle Database (2-Tier)

If you do not need to load the graph into the in-memory graph server, you can use another approach: create a property graph from data in relational tables, and store it in the property graph schema (VT\$ and GE\$ tables). You can then run PGQL queries on this graph.

You can load this graph into memory for running analytics algorithms and PGQL queries not supported in the database. You can configure the in-memory graph server to periodically fetch updates from the data automatically in the graph to keep the data synchronized.

This approach uses a two-tier architecture, as shown in the following figure.

Figure 1-3 Two-Tier Property Graph Architecture



1.2 Before You Begin with Oracle Property Graphs

Before you begin to create an Oracle Property Graph, you may need to adhere to one or more of the prerequisites explained in this section.

You must perform the tasks listed in [Table 1-2](#) to get started on property graphs.

Table 1-2 Overview of Tasks to Get Started with Property Graphs

Sequence	Task	Description	More Information
1	Verify Oracle Database Requirements	Ensure that your Oracle Database version is 12.2 and higher.	Database Compatibility and Restrictions
2	Download Oracle Graph Server and Client	Download Oracle Graph Server and Client from Oracle Software Delivery Cloud or from Oracle Technology Network .	Downloading Oracle Graph Server and Client
3	Install the PL/SQL patch in your Oracle Database	Upgrade the PL/SQL Graph packages in your Oracle Database.	Installing PL/SQL Packages in Oracle Database
4	Install Oracle Graph Server	Install Oracle Graph server, which is available as a separate downloadable package.	Installing Oracle Graph Server
5	Download Oracle Graph Client	Install Oracle Graph Client to work with property graphs.	Installing the Java Client
6	Set up transport layer security	Configure the graph server and client to trust the self-signed certificate.	Setting Up Transport Layer Security
7	Add permissions to publish the graph	Grant permissions to publish graphs.	Adding Permissions to Publish the Graph

- [Database Compatibility and Restrictions](#)
- [Downloading Oracle Graph Server and Client](#)
- [Installing PL/SQL Packages in Oracle Database](#)

Oracle Graph Server and Client will work with Oracle Database 12.2 onward. However, you must install the updated PL/SQL packages that are part of the Oracle Graph Server and Client download.

1.2.1 Database Compatibility and Restrictions

Oracle Graph Server and Client will work with Oracle Database 12.2 onward. This includes working with the family of Oracle Autonomous Database -- all versions of Oracle Autonomous Data Warehouse (shared), Oracle Autonomous Database (shared), and Oracle Autonomous Database (dedicated).

For details, including any limitations and actions you should take to address them, see "[Database Compatibility Matrix for Oracle Graph Server and Client](#)".

1.2.2 Downloading Oracle Graph Server and Client

You can download **Oracle Graph Server and Client** from [Oracle Software Delivery Cloud](#) or from [Oracle Technology Network](#).

Table 1-3 summarizes all the files contained in the Oracle Graph Server and Client deployment.

`<ver>` denoted in the file name in the **Table 1-3** reflects the downloaded Oracle Graph Server and Client version.

Table 1-3 Components in the Oracle Graph Server and Client Deployment

File	Component	Description
oracle-graph- <code><ver></code> .rpm	Oracle Graph Server	An rpm file to deploy Oracle Graph Server.
oracle-graph-client- <code><ver></code> .zip	Oracle Graph Client	A zip file containing Oracle Graph Client.
oracle-graph-zepplin-interpreter- <code><ver></code> .zip	Oracle Graph Apache Zeppelin Client	A zip file containing libraries to use Apache Zeppelin to work with Oracle Graph.
oracle-graph-hdfs-connector- <code><ver></code> .zip	Oracle Graph HDFS Connector	A zip file containing libraries to connect Oracle Graph Server with the Apache Hadoop Distributed Filesystem (HDFS).
oracle-graph-sqlcl-plugin- <code><ver></code> .zip	Oracle Graph PGQL Plugin for SQLcl	A plugin for SQLcl to run PGQL queries in SQLcl.
oracle-graph-webapps- <code><ver></code> .zip	Oracle Graph Web Applications	A zip file containing <code>.war</code> files for deploying graph servers in an application server.
oracle-graph-plsql- <code><ver></code> .zip	Oracle Graph PL/SQL Patch	A zip file containing PL/SQL packages. It is recommended to update the PL/SQL Graph packages in your database with these packages. Instructions are in the README file.

1.2.3 Installing PL/SQL Packages in Oracle Database

Oracle Graph Server and Client will work with Oracle Database 12.2 onward. However, you must install the updated PL/SQL packages that are part of the Oracle Graph Server and Client download.

Note:

You can skip this section if you are using Graph Server and Client with Oracle Autonomous Database. You only need to create roles and assign permissions by executing step-5 and step-6 in [Basic Steps for Using an Oracle Database for Authentication](#). You can run these steps using Database Actions in Oracle Cloud Infrastructure Console.

1. Download the Oracle Graph PL/SQL patch component, which is a part of the Oracle Graph Server and Client download from [Oracle Software Delivery Cloud](#).
2. Unzip the file `oracle-graph-plsql-<ver>.zip` into a directory of your choice.
`<ver>` denotes the version downloaded for the Oracle Graph PL/SQL Patch for PL/SQL.
3. Install the PL/SQL packages:
 - There are two directories, one for users with Oracle Database 18c or below, and one for users with Oracle Database 19c or above. As a database user with DBA privileges, follow the instructions in the README.md file in the appropriate directory (that matches your database version). This has to be done for every PDB you will use the graph feature in. For example:

```
-- Connect as SYSDBA
SQL> ALTER SESSION SET CONTAINER=<YOUR_PDB_NAME>;
SQL> @opgremov.sql
SQL> @catopg.sql
```

4. Create a database user in the database for working with graphs:
 - a. As a database user with DBA privileges, create a user `<graphuser>`, and grant the necessary privileges.
 - i. If you plan to use a three-tier architecture (graph queries and analytics executed in the in-memory graph server (PGX), then grant privileges as described in the following command:

```
SQL> GRANT CREATE SESSION, CREATE TABLE, CREATE VIEW TO <graphuser>
```

- ii. If you plan to use a two-tier architecture and run graph queries in the database, then grant privileges as described in [Required Privileges for Database Users](#):

```
SQL> GRANT CREATE SESSION, ALTER SESSION, CREATE TABLE, CREATE
PROCEDURE, CREATE TYPE, CREATE SEQUENCE, CREATE VIEW, CREATE
TRIGGER TO <graphuser>
```

- b. As a <graphuser> in the database, check that the PL/SQL update is successful:

```
SQL> CONNECT <graphuser>/<password>
SQL> SELECT opg_apis.get_opg_version() FROM DUAL;
-- Should return 21.4 if you are using
-- Graph Server and Client 21.4
```

5. Grant the appropriate roles (GRAPH_DEVELOPER or GRAPH_ADMINISTRATOR), to the database user created in step 4 for working with the graphs.

 **Note:**

- See [User Authentication and Authorization](#) for more information on authorization rules for Graph Server (PGX) and Client 21.4.
- See [Upgrading From Graph Server and Client 20.4.x to 21.x](#) for more information if you are migrating to Graph Server (PGX) and Client 21.1 from an earlier version.

```
SQL> GRANT GRAPH_DEVELOPER to <graphuser>
SQL> GRANT GRAPH_ADMINISTRATOR to <adminuser>
```

1.3 Oracle Graph Server and Client Installation

This section explains the various operations that you must perform to install, uninstall or upgrade Oracle Graph Server (PGX). It also includes the steps to install Oracle Graph Client.

- [Installing Oracle Graph Server](#)
- [Deploying Oracle Graph Server to a Web Server](#)
- [Upgrading Oracle Graph Server](#)
- [Uninstalling Oracle Graph Server](#)
- [Installing Oracle Graph Client](#)

1.3.1 Installing Oracle Graph Server

The prerequisites for installing the Oracle Graph Server are:

- Oracle Linux 6 , 7 or 8 x64 or a similar Linux distribution such as RedHat ([Using the Machine Learning Library \(PgxML\) for Graphs](#) requires Oracle Linux 7 or later)
- Oracle JDK 8 or JDK 11

 **Note:**

Due to a bug in Open JDK, it is recommended to avoid the following Oracle JDK versions:

- JDK 11.0.9
- JDK 11.0.10
- JDK 11.0.11
- JDK 11.0.12

See this [note](#) for more details.

You can run Oracle Graph Server in standalone mode or using a webserver like Oracle WebLogic Server or Apache Tomcat.

The installation steps for installing Oracle Graph Server in standalone mode are as shown:

1. As a `root` user or using `sudo`, install the RPM file using the `rpm` command line utility:

```
sudo rpm -i oracle-graph-<version>.rpm
```

Where `<version>` reflects the version that you downloaded. (For example: `oracle-graph-21.4.0.0.0.x86_64.rpm`)

The `.rpm` file is the graph server.

The following post-installation steps are carried out at the time of the RPM file installation:

- Creation of a working directory in `/opt/oracle/graph/pgx/tmp_data`
- Creation of a log directory in `/var/log/oracle/graph`
- Installation of Python Client

 **Note:**

If Python is not installed in your system, then this step will be skipped.

- Automatic generation of self-signed TLS certificates in `/etc/oracle/graph`

 **Note:**

- You can also choose to configure and set up transport layer security (TLS) in graph server. See [Setting Up Transport Layer Security](#) for more details.
- For demonstration purposes, if you wish to disable transport layer security (TLS) in graph server, see [Disabling Transport Layer Security \(TLS\) in Graph Server](#) for more details.

2. As `root` or using `sudo`, add operating system users allowed to use the server installation to the operating system group `oraclegraph`. For example:

```
usermod -a -G oraclegraph <graphuser>
```

This adds the specified graph user to the group `oraclegraph`.
Note that `<graphuser>` must log out and log in again for this to take effect.

3. As `<graphuser>`, configure the server by modifying the files under `/etc/oracle/graph` by following the steps under [Prepare the Graph Server for Database Authentication](#).
4. Ensure that authentication is enabled for database users that will connect to the graph server, as explained in [User Authentication and Authorization](#).
5. As a `root` user or using `sudo`, start the graph server (PGX) by executing the following command:

```
sudo systemctl start pgx
```

You can verify if the graph server has started by executing the following command:

```
systemctl status pgx
```

- If the graph server has successfully started, the response may appear as:
 - `pgx.service - Oracle Graph In-Memory Server`
Loaded: loaded (/etc/systemd/system/pgx.service; disabled; vendor preset: disabled)
Active: active (running) since Wed 2021-01-27 10:06:06 EST; 33s ago
Main PID: 32127 (bash)
CGroup: /system.slice/pgx.service
├─32127 /bin/bash start-server
└─32176 java -Dlog4j.configurationFile=/etc/oracle/graph/log4j2-server.xml -Doracle.jdbc.fanEnabled=false -cp /opt/oracle/graph/pgx/bin/../../pgx/server/lib/activat...

The graph server is now ready to accept requests.

- If the graph server has not started, then you must check the log files in `/var/log/oracle/graph` for errors. Additionally, you can also run the following command to view any `systemd` errors:

```
sudo journalctl -u pgx.service
```

Additional installation operations are required for specific use cases, such as:

- Analyze property graphs using Python (see [Installing the Python Client](#)).
- Deploy the graph server as a web application with Oracle WebLogic Server (see [Deploying to Oracle WebLogic Server](#)).
- Deploy GraphViz in Oracle WebLogic Server (see [Deploying the Graph Visualization Application in Oracle WebLogic Server](#)).

- Deploy the graph server as a web application with Apache Tomcat (see [Deploying to Apache Tomcat](#)).

For instructions to deploy the graph server in Oracle WebLogic Server or Apache Tomcat, see:

- [Deploying to Oracle WebLogic Server](#)
- [Deploying to Apache Tomcat](#)

1.3.2 Deploying Oracle Graph Server to a Web Server

You can deploy Oracle Graph Server to Apache Tomcat or Oracle WebLogic Server.

The following explains the deployment instructions:

- [Deploying to Apache Tomcat](#)
The example in this topic shows how to deploy the graph server as a web application with Apache Tomcat.
- [Deploying to Oracle WebLogic Server](#)
The example in this topic shows how to deploy the graph server as a web application with Oracle WebLogic Server.

1.3.2.1 Deploying to Apache Tomcat

The example in this topic shows how to deploy the graph server as a web application with Apache Tomcat.

The graph server will work with Apache Tomcat 9.0.x.

1. Download the Oracle Graph Webapps zip file from [Oracle Software Delivery Cloud](#). This file contains ready-to-deploy Java web application archives (.war files). The file name will be similar to this: `oracle-graph-webapps-<version>.zip`.
2. Unzip the file into a directory of your choice.
3. Locate the .war file that follows the naming pattern: `graph-server-<version>-pgx<version>.war`.
4. Configure the graph server.
 - a. Modify authentication and other server settings by modifying the `WEB-INF/classes/pgx.conf` file inside the web application archive. See [User Authentication and Authorization](#) section for more information.
 - b. Optionally, change logging settings by modifying the `WEB-INF/classes/log4j2.xml` file inside the web application archive.
 - c. Optionally, change other servlet specific deployment descriptors by modifying the `WEB-INF/web.xml` file inside the web application archive.
5. Copy the .war file into the Tomcat webapps directory. For example:

```
cp graph-server-<version>-pgx<version>.war $CATALINA_HOME/webapps/pgx.war
```

 **Note:**

The name you give the war file in the Tomcat `webapps` directory determines the context path of the graph server application. It is recommended naming the war file as `pgx.war`.

6. Configure Tomcat specific settings, like the correct use of TLS/encryption.
7. Ensure that port 8080 is not already in use.
8. Start Tomcat:

```
cd $CATALINA_HOME
./bin/startup.sh
```

The graph server will now listen on `localhost:8080/pgx`.

You can connect to the server from JShell by running the following command:

```
$ <client_install_dir>/bin/opg4j --base_url https://
localhost:8080/pgx -u <graphuser>
```

Related Topics

- [The Tomcat documentation \(select desired version\)](#)

1.3.2.2 Deploying to Oracle WebLogic Server

The example in this topic shows how to deploy the graph server as a web application with Oracle WebLogic Server.

This example shows how to deploy the graph server with Oracle WebLogic Server. Graph server supports WebLogic Server version 12.1.x and 12.2.x.

1. Download the Oracle Graph Webapps zip file from [Oracle Software Delivery Cloud](#). This file contains ready-to-deploy Java web application archives (`.war` files). The file name will be similar to this: `oracle-graph-webapps-<version>.zip`.
2. Unzip the file into a directory of your choice.
3. Locate the `.war` file that follows the naming pattern: `graph-server-<version>-pgx<version>.war`.
4. Configure the graph server.
 - a. Modify authentication and other server settings by modifying the `WEB-INF/classes/pgx.conf` file inside the web application archive.
 - b. Optionally, change logging settings by modifying the `WEB-INF/classes/log4j2.xml` file inside the web application archive.
 - c. Optionally, change other servlet specific deployment descriptors by modifying the `WEB-INF/web.xml` file inside the web application archive.
 - d. Optionally, change WebLogic Server-specific deployment descriptors by modifying the `WEB-INF/weblogic.xml` file inside the web application archive.

5. Configure WebLogic specific settings, like the correct use of TLS/encryption.
6. Deploy the `.war` file to WebLogic Server. The following example shows how to do this from the command line:

```
. $MW_HOME/user_projects/domains/mydomain/bin/setDomainEnv.sh
. $MW_HOME/wlserver/server/bin/setWLSEnv.sh
java weblogic.Deployer -adminurl http://localhost:7001 -username
<username> -password <password> -deploy -source <path-to-war-file>
```

- [Installing Oracle WebLogic Server](#)

1.3.2.2.1 Installing Oracle WebLogic Server

To download and install the latest version of Oracle WebLogic Server, see

<http://www.oracle.com/technetwork/middleware/weblogic/documentation/index.html>

1.3.3 Upgrading Oracle Graph Server

To upgrade the graph server, make sure the graph server is shut down, then execute the following command with the newer RPM file as an argument.

- Run the following command as a `root` user or with `sudo`:

```
sudo rpm -U oracle-graph-21.4.0.0.0.x86_64.rpm
```

1.3.4 Uninstalling Oracle Graph Server

To uninstall the graph server, make sure the graph server is shut down.

- Run the following command as a `root` user or with `sudo`:

```
sudo rpm -e oracle-graph
```

1.3.5 Installing Oracle Graph Client

This sections explains in detail the installation steps for the various clients.

- [Installing the Java Client](#)
- [Installing the Python Client](#)
- [Uninstalling the Python Client](#)
This section describes how to uninstall the Python client.
- [Enabling the Graph Visualization Application](#)
- [Deploying the Graph Visualization Application](#)
This section describes the various methods to deploy the Graph Visualization Application.
- [Installing the Graph Zeppelin Interpreter Client](#)

1.3.5.1 Installing the Java Client

The prerequisites for installing the Java client are:

- A Unix-based operation system (such as Linux) or macOS or Microsoft Windows
- Oracle JDK 11

 **Note:**

Due to a bug in Open JDK, which causes a deadlock when you attempt to copy and paste into a JShell session, it is recommended that you avoid the following Oracle JDK versions:

- JDK 11.0.9
- JDK 11.0.10
- JDK 11.0.11
- JDK 11.0.12

1. Download Oracle Graph Client 21.4 from [Oracle Software Cloud](#).
2. Unzip the file into a directory of your choice.
3. Configure your client to trust the self-signed server certificate. See [Configuring a Client to Trust the Self-Signed Certificate](#) for more information.
4. Start the OPG4J shell to connect to the graph server (PGX) as shown:

```
cd <CLIENT_INSTALL_DIR>
./bin/opg4j --base_url https://<host>:7007 --username <graphuser>
```

In the preceding code:

- **<CLIENT_INSTALL_DIR>**: Directory where the shell executables are located. The shell executables are generally found in `/opt/oracle/graph/bin` after server installation, and `<CLIENT_INSTALL_DIR>/bin` after the client installation.
- **<host>**: Server host

 **Note:**

The graph server (PGX), listens on port `7007` by default. If needed, you can configure the graph server to listen on a different port by changing the port value in the server configuration file (`server.conf`). See [Configuring the In-Memory Graph Server \(PGX\)](#) for details.

- **<graphuser>**: Database user

You will be prompted for the database password.

See [Starting the OPG4J Shell](#) for more information on the different ways you can start the OPG4J shell.

The OPG4J shell starts and the following command line prompt appears as shown:

```
For an introduction type: /help intro
Oracle Graph Server Shell 21.2.0
Variables instance, session, and analyst ready to use.
opg4j>
```

1.3.5.2 Installing the Python Client

To install the Python client, you must ensure that your system meets the prerequisites mentioned in [Prerequisites for Installing the Python Client](#).

You can execute the following steps to install and connect using the Python client:

1. Download the Oracle Graph Client from [Oracle Software Cloud](#).

For example, `oracle-graph-client-21.4.0.zip`.

2. Install the client through `pip`.

For example,

```
pip3 install --user oracle-graph-client-21.4.0.zip
```

3. Configure your client to trust the self-signed server certificate. See [Configuring a Client to Trust the Self-Signed Certificate](#) for more information.
4. Start the OPG4Py shell to connect to the graph server(PGX) by running the following command:

```
cd <CLIENT_INSTALL_DIR>
./bin/opg4py --base_url https://<host>:7007
```

In the preceding code:

- **<CLIENT_INSTALL_DIR>**: Directory where the shell executables are located. The shell executables are generally found in `/opt/oracle/graph/bin` after server installation, and `<CLIENT_INSTALL_DIR>/bin` after the client installation.
- **<host>**: Server host

Note:

The graph server (PGX), listens on port 7007 by default. If needed, you can configure the graph server to listen on a different port by changing the port value in the server configuration file (`server.conf`). See [Configuring the In-Memory Graph Server \(PGX\)](#) for details.

You are prompted to enter your username and password.

See [Starting the OPG4Py Shell](#) for more information on the different ways you can start the OPG4Py shell.

The OPG4Py shell starts and the following command line prompt appears as shown:

```
Oracle Graph Server Shell 21.4.0  
>>>
```

 **Note:**

You can also install the python client library in Jupyter Notebook. Using the Python API, you can then connect to the graph server (PGX) to run PGQL queries and graph algorithms in a Jupyter Notebook environment.

- [Prerequisites for Installing the Python Client](#)

1.3.5.2.1 Prerequisites for Installing the Python Client

You must ensure that the following prerequisites are met before you install the Python client:

1. Make sure that the following softwares are installed in your system:
 - Oracle JDK 8 or later
 - Python 3.6 or later
To verify you are using the right version of the Python client, run the following command:

```
$> python3 --version  
Python 3.6.1
```

 **Note:**

Python 2.x is not supported.
For more information on installing Python 3 on Oracle Linux, see [Python for Oracle Linux](#).

2. Ensure that `python3-devel` is installed in your system.

```
sudo yum install python3-devel
```

1.3.5.3 Uninstalling the Python Client

This section describes how to uninstall the Python client.

To uninstall the Python client, run the following command:

```
pip3 uninstall pypgx
```

1.3.5.4 Enabling the Graph Visualization Application

There are two ways you can use the Graph Visualization application:

- **Standalone mode**
If you install the Graph Server rpm file, the Graph Visualization application starts up by default when you start the PGX server.
- **Custom web container mode**
You can download the `oracle-graph-webapps-<version>.zip` package which contains a web application archive (WAR) file. You can deploy this file into your Oracle Weblogic 12.2 (or later) or Apache Tomcat (9.x or later) web containers.

See [Deploying the Graph Visualization Application](#) for more information.

The Graph Visualization application requires the Oracle Graph Server to be installed as a prerequisite component.

See [Installing Oracle Graph Server](#) for more information.

To start the Graph Visualization application in standalone mode:

1. Start the graph server (PGX) as shown:

```
sudo systemctl start pgx
```

The Graph Visualization application starts up by default.

2. Configure your Graph Visualization application to trust the self-signed server certificate. See [Configuring a Client to Trust the Self-Signed Certificate](#) for more information.
3. Connect to your browser for running the Graph Visualization application as shown

```
https://localhost:7007/ui
```

One of the following messages may appear:

- Your connection is not private
- Your connection is not secure

Click the Continue or Accept button to proceed.

1.3.5.5 Deploying the Graph Visualization Application

This section describes the various methods to deploy the Graph Visualization Application.

- [Deploying the Graph Visualization Application in Standalone Mode](#)
You can use the instructions in this section to deploy the Graph Visualization application in a standalone mode.
- [Deploying the Graph Visualization Application to Apache Tomcat](#)
- [Deploying the Graph Visualization Application in Oracle WebLogic Server](#)
The following instructions are for deploying the Graph Visualization application in Oracle WebLogic Server 12.2.1.3. You might need to make slight modifications, as appropriate, for different versions of the Weblogic Server.

- [Configuring Advanced Options for PGQL Driver Selection](#)


1.3.5.5.1 Deploying the Graph Visualization Application in Standalone Mode

You can use the instructions in this section to deploy the Graph Visualization application in a standalone mode.

The Graph Visualization application starts up by default when you start the graph server (PGX server).

- Navigate to `https://localhost:7007/ui` in your browser. The Graph Visualization Login screen opens as shown:

Figure 1-4 Graph Visualization Login



ORACLE®
Graph Visualization

Username

Password

► Advanced Options

SUBMIT

- Enter your database **Username** and **Password**.
- Select and configure the required PGQL Driver.
See [Configuring Advanced Options for PGQL Driver Selection](#) for more information.
- Click **Submit**.
You are now signed into the Graph Visualization application.

The title bar on the query visualization page displays the connection mode along with the relevant URL.

1.3.5.5.2 Deploying the Graph Visualization Application to Apache Tomcat

The following are the steps to deploy the Graph Visualization application to Apache Tomcat.

1. Download the Oracle Graph Webapps zip file from [Oracle Software Delivery Cloud](#). This file contains ready-to-deploy Java web application archives (.war files). The file name will be similar to this: `oracle-graph-webapps-<version>.zip`.
2. Configure Tomcat specific settings, like the correct use of TLS/encryption.
3. Ensure that port 8080 is not already in use.
4. Start Tomcat:

```
cd $CATALINA_HOME
./bin/startup.sh
```

The Graph Visualization application is now listening on `localhost:8080/ui`

Note:

The name you give the war file (`graphviz-<version>-pgviz<graphviz-version>-tomcat.war`) in the Tomcat webapps directory determines the context path of the graph server application. It is recommended naming the war file as `ui.war`.

5. Navigate to the Graph Visualization Application using the URL, `localhost:8080/ui` in your browser.

The Graph Visualization login page appears as shown in [Figure 1-4](#).

6. Enter your database credentials and configure the required PGQL driver.
See [Configuring Advanced Options for PGQL Driver Selection](#) for more information.
7. Click Submit.

You are now signed into the Graph Visualization application.

The title bar on the query visualization page displays the connection mode along with the relevant URL.

1.3.5.5.3 Deploying the Graph Visualization Application in Oracle WebLogic Server

The following instructions are for deploying the Graph Visualization application in Oracle WebLogic Server 12.2.1.3. You might need to make slight modifications, as appropriate, for different versions of the Weblogic Server.

1. Download the Oracle Graph Webapps zip file from [Oracle Software Delivery Cloud](#). This file contains ready-to-deploy Java web application archives (.war files). The file name will be similar to this: oracle-graph-webapps-<version>.zip
2. Start WebLogic Server.

```
# Start Server
cd $MW_HOME/user_projects/domains/base_domain
./bin/startWebLogic.sh
```

3. Enable tunneling.
In order to be able to deploy the Graph Visualization application WAR file over HTTP, you must enable tunneling first. Go to the WebLogic admin console (by default on <http://localhost:7001/console>). Select **Environment** (left panel) > **Servers** (left panel). Click the server that will run Graph Visualization (main panel). Select (top tab bar), check **Enable Tunneling**, and click **Save**.
4. Deploy the graphviz-<version>-pgviz<graphviz-version>-wls.war file.
To deploy the WAR file to WebLogic Server, use the following command, replacing the <<...>> markers with values matching your installation:

```
cd $MW_HOME/user_projects/domains/base_domain
source bin/setDomainEnv.sh
java weblogic.Deployer -adminurl <<admin-console-url>> -username
<<admin-user>> -password <<admin-password>> -deploy -upload <<path/
to>>/graphviz-<<version>>-pgviz<<graphviz-version>>.war
```

To undeploy, you can use the following command:

```
java weblogic.Deployer -adminurl <<admin-console-url>> -username
<<admin-user>> -password <<admin-password>> -name <<path/to>>/
graphviz-<<version>>-pgviz<<graphviz-version>>.war -undeploy
```

To test the deployment, navigate using your browser to: <https://<<fqdn-ip>>:<<port>>/ui>.

The Graph Visualization Login screen appears as shown in [Figure 1-4](#).

5. Enter your database credentials and configure the required PGQL driver. See [Configuring Advanced Options for PGQL Driver Selection](#) for more information.
6. Click Submit.
You are now logged in and the Graph Visualization query user interface (UI) appears and the graphs from PGX are retrieved.

The title bar on the query visualization page displays the connection mode along with the relevant URL.

1.3.5.5.4 Configuring Advanced Options for PGQL Driver Selection

The Graph Visualization application can be configured to communicate either with the graph server (PGX) or to the Oracle Database. You can apply the required configuration at the time of login through the **Advanced Options** settings in the Graph Visualization login page.

You can dynamically change and configure the PGQL driver by following the instructions as appropriate for your preference:


- [Configuring the Graph Visualization Application for PGQL on Graph Server \(PGX\)](#)
- [Configuring the Graph Visualization Application for PGQL on Database](#)

1.3.5.5.4.1 Configuring the Graph Visualization Application for PGQL on Graph Server (PGX)

To configure Graph Visualization application to communicate with a PGX deployment (PGQL on Graph Server):

1. Click **Advanced Options** in the Graph Visualization login page.
2. Select **Graph Server** as shown:

Figure 1-5 PGQL on Graph Server (PGX)



3. Optionally, modify your **PGX Base URL**.

 **Note:**

- By default, the Graph Visualization application connects to the graph server (PGX) using the PGX base URL defined in the `web.xml` file for your installation.
- If you wish to disable transport layer security (TLS) in graph server, see [Disabling Transport Layer Security \(TLS\) in Graph Server](#) for more details.

4. Optionally, enter **Session Id**.

When the Graph Visualization application is using PGQL on Graph Server (PGX), the application will use your Oracle Database as identity manager by default. This means that you log into the application using existing Oracle Database credentials (username and password), and the actions which you are allowed to do on the graph server are determined by the roles that have been granted to you in the Oracle Database.

 **Note:**

If you wish to enable Kerberos Authentication for the Graph Visualization Application, see [Kerberos Enabled Authentication for the Graph Visualization Application](#) for more information.


1.3.5.5.4.2 Configuring the Graph Visualization Application for PGQL on Database

To configure the Graph Visualization application to communicate with Oracle Database (PGQL on Database):

1. Click **Advanced Options** in the Graph Visualization login page.
2. Select **Database** as shown:

Figure 1-6 PGQL on Database

▲ Advanced Options



3. Optionally, modify the **JDBC URL** for your Oracle database.

 **Note:**

- By default, the Graph Visualization application connects to the database using the JDBC URL defined in the `web.xml` file for your installation.
- If you wish to enable Kerberos Authentication for the Graph Visualization Application, see [Kerberos Enabled Authentication for the Graph Visualization Application](#) for more information.
- If you wish to disable transport layer security (TLS) in graph server, see [Disabling Transport Layer Security \(TLS\) in Graph Server](#) for more details.

1.3.5.6 Installing the Graph Zeppelin Interpreter Client

To install the graph interpreter into your local Zeppelin installation:

 **Note:**

The following steps were tested with Zeppelin version 0.9, and might have to be modified with newer versions.

As a prerequisite, you must have Java 8 set in your system for installing Apache Zeppelin.

1. Download and install [Apache Zeppelin](#).
2. Download and install [Apache Groovy 2.4.x](#)
3. Copy the following libraries:
 - a. Copy the libraries from the Oracle Graph Client for Apache Zeppelin package into `$ZEPPELIN_HOME/interpreter/pgx`.

```
unzip oracle-graph-zeppelin-interpreter-21.1.0.zip -d $ZEPPELIN_HOME/interpreter/pgx
```

- b. Copy the libraries inside `$GROOVY_HOME/lib` into `$ZEPPELIN_HOME/interpreter/pgx`.

```
cp $GROOVY_HOME/lib/* $ZEPPELIN_HOME/interpreter/pgx
```

4. Configure your graph Zeppelin interpreter client application to trust the self-signed server certificate. See [Configuring a Client to Trust the Self-Signed Certificate](#) for more information.
5. Restart Zeppelin.

1.4 Setting Up Transport Layer Security

The graph server (PGX), by default, allows only encrypted connections using Transport Layer Security (TLS). TLS requires the server to present a server certificate to the client and the client must be configured to trust the issuer of that certificate.

Starting with Graph Server and Client Release 21.1, the RPM file installation generates a self-signed certificate into `/etc/oracle/graph`, which the server uses to enable TLS by default. If self-signed certificates are sufficient for you to get started and if your connections are only to `localhost`, you can skip to [Configuring a Client to Trust the Self-Signed Certificate](#).

- [Generating a Self-Signed Server Certificate](#)
You can create a self-signed server certificate using the `openssl` command.
- [Configuring the Graph Server \(PGX\)](#)
You must specify the path to the server certificate and the server's private key in PEM format in the graph server (PGX) configuration file.
- [Configuring a Client to Trust the Self-Signed Certificate](#)
You must configure your client application to accept the self-signed graph server (PGX) certificate.

1.4.1 Generating a Self-Signed Server Certificate

You can create a self-signed server certificate using the `openssl` command.

The following steps show how to generate a self-signed server certificate.

1. Go to the following directory:

```
cd /etc/oracle/graph
```

2. Execute the following commands:

```
openssl req -new -newkey rsa:2048 -days 365 -nodes -x509 -subj "/C=US/ST=MyState/L=MyTown/O=MyOrganization/CN=ROOT" -keyout ca_key.pem -out ca_certificate.pem
openssl genrsa -out server_key_traditional.pem 2048
openssl pkcs8 -topk8 -in server_key_traditional.pem -inform pem -out server_key.pem -outform pem -nocrypt
openssl req -new -subj "/C=US/ST=MyState/L=MyTown/O=MyOrganization/CN=localhost" -key server_key.pem -out server.csr
chmod 600 server_key.pem
openssl x509 -req -CA ca_certificate.pem -CAkey ca_key.pem -in server.csr -out server_certificate.pem -days 365 -CAcreateserial
chown oraclegraph:oraclegraph server_key.pem
```

Note:

- The certificate mentioned in the above example will only work for the host `localhost`. If you have a different domain, you must replace `localhost` with your domain name.
- The above self-signed certificate is valid only for 365 days.

1.4.2 Configuring the Graph Server (PGX)

You must specify the path to the server certificate and the server's private key in PEM format in the graph server (PGX) configuration file.

Note:

If you deploy the graph server into your web server using the web applications download package, then this section does not apply. Please refer to the manual of your web server for instructions on how to configure TLS.

1. Edit the file at `/etc/oracle/graph/server.conf`, and specify the paths to the server certificate and the server's private key in PEM format, as shown:

```
{
  "port": 7007,
  "enable_tls": true,
  "server_private_key": "/etc/oracle/graph/server_key.pem",
  "server_cert": "/etc/oracle/graph/server_certificate.pem",
  "enable_client_authentication": false,
  "working_dir": "/opt/oracle/graph/pgx/tmp_data"
}
```

2. Restart the graph server.

 **Note:**

- You should use a certificate issued by a certificate authority (CA) which is trusted by your organization. If you do not have a CA certificate, you can temporarily create a self-signed certificate and get started.
- Always use a valid certificate trusted by your organization. We do not recommend the usage of self-signed certificates for production environments.

1.4.3 Configuring a Client to Trust the Self-Signed Certificate

You must configure your client application to accept the self-signed graph server (PGX) certificate.

To configure a client to trust the self-signed certificate, the root certificate must be imported to your Java installation local trust store.

- For a Java or a Python client, you must import the root certificate to all the Java installations used by all the clients.

 **Note:**

The JShell client requires Java 11.

- For the Graph Visualization application, you must import the root certificate to the system Java installation of the environment running the graph server (PGX) or the web server serving the graph visualization application. That is, the JDK installation which is used by the OS user running the server that serves the Graph Visualization application.
- For the Graph Zeppelin interpreter client, you must import the root certificate to the Java installation used by the Zeppelin server.

You can import the root certificate as shown in the following step:

- Execute the following command as a `root` user or with `sudo`:

1. For Java 8 (make sure `JAVA_HOME` is set):

```
sudo keytool -import -trustcacerts -keystore $JAVA_HOME/jre/lib/  
security/cacerts -storepass changeit -alias pgx -file /etc/oracle/  
graph/ca_certificate.pem -noprompt
```

2. For Java 11 (make sure `JAVA11_HOME` is set):

```
sudo keytool -import -trustcacerts -keystore $JAVA11_HOME/lib/  
security/cacerts -storepass changeit -alias pgx -file /etc/oracle/  
graph/ca_certificate.pem -noprompt
```

where `changeit` is the sample keystore password. You can change this password to a password of your choice. Be sure to remember this password as you will need it to modify the certificate.

1. If you are upgrading the graph server from a previous release, you must first delete the certificate by executing the following command appropriate to your Java version. You must run the command using `sudo` or as a `root` user:

For Java 8:

```
sudo keytool -delete -alias pgx -keystore $JAVA_HOME/jre/lib/  
security/cacerts -storepass changeit
```

For Java 11:

```
sudo keytool -delete -alias pgx -keystore $JAVA11_HOME/lib/  
security/cacerts -storepass changeit
```

2. Import the new certificate as shown in the preceding [step](#).

1.5 Adding Permissions to Publish the Graph

There are two ways by which you can view any graph in your graph server (PGX) session in the graph visualization application.

When you log into the graph visualization tool in your browser, that will be a different session from your JShell session or application session. To visualize the graph you are working on in your JShell session or application session in your graph visualization session, you can perform one of the following two steps:

1. Get the session id of your working session using the `PgxSession` API, and use that session id when you log into the graph visualization application. This is the recommended option.

```
opg4j> session.getId();  
$2 ==> "898bdbc3-af80-49b7-9a5e-10ace6c9071c" //session id
```

or

2. Grant `PGX_SESSION_ADD_PUBLISHED_GRAPH` permission and then publish the graph as shown:

- a. Grant `PGX_SESSION_ADD_PUBLISHED_GRAPH` role in the database to the user visualizing the graph as shown in the following statement:

```
GRANT PGX_SESSION_ADD_PUBLISHED_GRAPH TO <graphuser>
```

- b. Publish the graph when you are ready to visualize the graph using the `publish` API.

 **Note:**

- See [User Authentication and Authorization](#) for more information on authorization rules for Graph Server (PGX) and Client 21.1.
- See [Upgrading From Graph Server and Client 20.4.x to 21.x](#) for more information if you are migrating to Graph Server (PGX) and Client 21.4 from an earlier version.

1.6 Security Best Practices with Graph Data

Several security-related best practices apply when working with graph data.

Sensitive Information

Graph data can contain sensitive information and should therefore be treated with the same care as any other type of data. Oracle recommends the following considerations when using a graph product:

- Avoid storing sensitive information in your graph if that information is not required for analysis. If you have existing data, only model the relevant subset you need for analysis as a graph, either by applying a preprocessing step or by using subgraph and filtering techniques that are part of graph product.
- Model your graph in a way that vertex and edge identifiers are not considered sensitive information.
- Do not deploy the product into untrusted environments or in a way that gives access to untrusted client connections.
- Make sure all communication channels are encrypted and that authentication is always enabled, even if running within a trusted network.

Least Privilege Accounts

The database user account that is being used by the in-memory analyst (PGX) to read data should be a low-privilege, read-only account. PGX is an in-memory accelerator that acts as a read-only cache on top of the database, and it does not write any data back to the database.

If your application requires writing graph data and later analyzing it using PGX, make sure you use two different database user accounts for each component.

1.7 Interactive Graph Shell CLIs

Both the Oracle Graph server and client packages contain interactive command-line applications for interacting with the Java APIs and the Python APIs of the product, locally or on remote computers.

The interactive graph shells dynamically interpret command-line inputs from the user, execute them by invoking the underlying functionality, and can print results or process them further. The graph shells provide a lightweight and interactive way of exercising graph functionality without creating a Java or Python application.

The graph shells are especially helpful if you want to do any of the following:

- Quickly run a "one-off" graph analysis on a specific data set, rather than creating a large application
- Run getting started examples and create demos on a sample data set
- Explore the data set, trying different graph analyses on the data set interactively
- Learn how to use the product and develop a sense of what the built-in algorithms are good for
- Develop and test custom graph analytics algorithms

The graph shell for the Java API (OPG4J) is implemented on top of the Java Shell tool (JShell). As such, it inherits all features provided by JShell such as tab-completion, history, reverse search, semicolon inference, script files, and internal variables. The graph shell for the Python API (OPG4Py) uses IPython in case it is installed.

The following sections explain in detail on how to start the graph shell CLIs:

- [Starting the OPG4J Shell](#)
- [Starting the OPG4Py Shell](#)

1.7.1 Starting the OPG4J Shell

Launching the OPG4J Shell

The Java shell executables are found in `/opt/oracle/graph/bin` after the graph server (PGX) installation, and in `<CLIENT_INSTALL_DIR>/bin` after the Java client installation.

The OPG4J shell uses JShell, which means the shell needs to run on Java 11 or later. See [Installing the Java Client](#) for more details on the prerequisites. You can then launch the OPG4J shell by entering the following in your terminal:

```
cd /opt/oracle/graph
./bin/opg4j
```

When the shell has started, the following command line prompt appears:

```
For an introduction type: /help intro
Oracle Graph Server Shell 21.4.0
Variables instance, session, and analyst ready to use.
opg4j>
```

By default, the OPG4J shell creates a local PGX instance, to run graph functions in the same JVM as the shell as described in [Developing Applications Using Graph Server Functionality as a Library](#).

Command-line Options

To view the list of available command-line options, add `--help` to the `opg4j` command:

```
./bin/opg4j --help
```

To start the `opg4j` shell without connecting to the graph server (PGX), use the `--no_connect` option as shown:

```
./bin/opg4j --no_connect
```

Starting the OPG4J Shell on Remote Mode

The OPG4J shell can connect to a graph server (PGX) instance that is running on another JVM (possibly on a different machine). In order to launch the OPG4J shell in remote mode, you must specify the `--base_url` parameter as shown:

```
./bin/opg4j --base_url https://<host>:7007 --username <graphuser>
```

where :

- `<host>`: is the server host
- `<graphuser>`: is the database user
You will be prompted for the database password.

Note:

The graph server (PGX), listens on port `7007` by default. If needed, you can configure the graph server to listen on a different port by changing the port value in the server configuration file (`server.conf`). See [Configuring the In-Memory Graph Server \(PGX\)](#) for details.

When the shell has started, the following command line prompt appears:

```
Oracle Graph Server Shell 21.4.0  
Variables instance, session, and analyst ready to use.  
opg4j>
```

If you have multiple versions of Java installed, you can easily switch between installations by setting the `JAVA_HOME` variable before starting the shell. For example:

```
export JAVA_HOME=/usr/lib/jvm/java-11-oracle
```

Batch Execution of Scripts

The OPG4J shell can execute a script by passing the path(s) to the script(s) to the `opg4j` command. For example:

```
./bin/opg4j /path/to/script.jsh
```

Predefined Functions

The OPG4J shell provides the following utility functions:

- `println(String)`: A shorthand for `System.out.println(String)`.

- `loglevel(String loggerName, String levelName)`: A convenient function to set the loglevel.

The `loglevel` function allows you to set the log level for a logger. For example, `loglevel("ROOT", "INFO")` sets the level of the root logger to `INFO`. This causes all logs of `INFO` and higher (`WARN`, `ERROR`, `FATAL`) to be printed to the console.

Script Arguments

You can also provide parameters to the script. For example:

```
./bin/opg4j /path/to/script.jsh script-arg-1 script-arg-2
```

In this example, the script `/path/to/script.jsh` can access the arguments via the `scriptArgs` system property. For example:

```
println(System.getProperty("scriptArgs"))// Prints: script-arg-1  
script-arg-2
```

Staying in Interactive Mode

By default, the OPG4J shell exits after it finishes execution. To stay in interactive mode after the script finishes *successfully*, pass the `--keep_running` flag to the shell. For example:

```
./bin/opg4j -b https://myserver.com:7007/ /path/to/script.jsh --  
keep_running
```

1.7.2 Starting the OPG4Py Shell

Launching the OPG4Py Shell

The OPG4Py shell executables are found in `/opt/oracle/graph/bin` after the graph server (PGX) installation, and in `<CLIENT_INSTALL_DIR>/bin` after the Python client installation.

Before launching the OPG4Py shell, verify that your system meets the prerequisites explained in [Prerequisites for Installing the Python Client](#). You can then launch the OPG4Py shell by entering the following in your terminal:

```
cd /opt/oracle/graph  
./bin/opg4py
```

When the shell has started, the following command line prompt appears:

```
Oracle Graph Server Shell 21.4.0  
>>>
```

If IPython is installed the following prompt will appear:

```
In [1]:
```

By default, the OPG4Py shell creates a local PGX instance, to run graph functions in the same JVM as the shell as described in [Developing Applications Using Graph Server Functionality as a Library](#).

Command-line Options

To view the list of available command-line options, add `--help` to the `opg4py` command:

```
./bin/opg4py --help
```

To start the PyPGX shell without connecting to the graph server (PGX), use the `--no_connect` option as shown:

```
./bin/opg4py --no_connect
```

Starting the OPG4Py Shell on Remote Mode

The OPG4Py shell can connect to a graph server (PGX) instance that is running on another JVM (possibly on a different machine). In order to launch the OPG4Py shell in remote mode, you must specify the `--base_url` parameter as shown:

```
./bin/opg4py --base_url https://<host>:7007 --username <graphuser>
```

where :

- `<host>`: is the server host
- `<graphuser>`: is the database user
You will be prompted for the database password.



Note:

The graph server (PGX), listens on port 7007 by default. If needed, you can configure the graph server to listen on a different port by changing the port value in the server configuration file (`server.conf`). See [Configuring the In-Memory Graph Server \(PGX\)](#) for details.

When the OPG4Py shell has started, the following command line prompt appears:

```
Oracle Graph Server Shell 21.4.0  
>>>
```

1.8 Developing Applications Using Graph Server Functionality as a Library

The graph functions available with the graph server (PGX) can be used as a library in your application.

After the rpm install of the graph server, all the jar files can be found in `/opt/oracle/graph/lib`. In this case, the server installation and the client user application are in the same machine.

For such use cases, development and testing can be done using the interactive Java shell or the Python shell in embedded (local) mode. This means a local PGX instance is created and runs in the same JVM as the client. If you start the shell without any parameters it will start a local PGX instance and run in embedded mode.

See [Using Graph Server \(PGX\) as a Library](#) for more information to obtain reference to a local PGX instance.

1.9 Storing Graphs in Oracle Database and Loading Graphs into Memory

You can work with graphs in **two-tier mode** (graph client connects directly to Oracle Database), or **three-tier mode** (graph client connects to the graph server (PGX) on the middle-tier, which then connects to Oracle Database).

Both modes for connecting to Oracle Database can be used regardless of whether the database is autonomous or not autonomous.

The database schema storing the graph must have the privileges listed in [Required Privileges for Database Users](#).

If you are using the Oracle Autonomous Database, see also [Using Oracle Graph with the Autonomous Database](#) for information about two-tier and three-tier deployments.

- [Two-Tier Mode](#)
In two-tier mode, the client graph application connects directly to Oracle Database.
- [Three-Tier Mode](#)
In three-tier mode, the client graph application connects to the graph server (PGX) in the middle tier, and the graph server connects to Oracle Database.

1.9.1 Two-Tier Mode

In two-tier mode, the client graph application connects directly to Oracle Database.

The graph is stored in the property graph schema (see [Property Graph Schema Objects for Oracle Database](#)).

You can use the PGQL DDL statement [CREATE PROPERTY GRAPH](#) to create a graph from database tables and store it in the property graph schema. You can then run PGQL queries on this graph from JShell shell, Java application, or the graph visualization tool.

The graph can be loaded from the property graph schema into memory in the graph server for faster processing and for using the analytics API.

1.9.2 Three-Tier Mode

In three-tier mode, the client graph application connects to the graph server (PGX) in the middle tier, and the graph server connects to Oracle Database.

The graph can be loaded from the property graph schema into the graph server, or directly from database tables into the graph server.

- **Loading a Graph from Property Graph Schema:**

Loading a graph from the property graph schema into memory in the graph server is the same as in the two-tier mode.

- **Loading a Graph Directly from Database Tables:**

When you load the graph from database tables into memory in the graph server, you create the graph in memory by directly reading data from the database tables. You do not create a graph in the property graph schema.

For more information about loading a graph from database tables into memory, see [Store the Database Password in a Keystore](#).

After the graph is loaded into memory, you can run PGQL queries on this graph from JShell shell, Java application, or the graph visualization tool. You can run graph analytics API from JShell shell or Java application, and visualize the results in the graph visualization application (GraphViz).

1.10 Using Oracle Graph with the Autonomous Database

Oracle Graph with the Autonomous Database allows you to create property graphs from data in your Autonomous Database.

When using Oracle Autonomous Database in a shared deployment, you can use Graph Studio, a powerful user interface for developing applications that use graph analysis. Using Graph Studio, you can automate the modeling of graphs from tables in Autonomous Database. You can interactively analyze and visualize the graph queries using advanced notebooks with multiple visualization options. You can execute over 60 built-in graph algorithms in Graph Studio to gain useful insights on your graph data. See [Using Graph Studio in Oracle Autonomous Database](#) for more information.

Alternatively, you can use Oracle Graph Server and Client with the family of Oracle Autonomous Database to create and work with property graphs.

This includes all versions of Oracle Autonomous Data Warehouse (shared), Oracle Autonomous Database (shared), and Oracle Autonomous Database (dedicated).

You can connect in two-tier mode (connect directly to Autonomous Database) or three-tier mode (connect to PGX on the middle tier, which then connects to Autonomous Database). (For basic information about two-tier and three-tier connection modes, see [Storing Graphs in Oracle Database and Loading Graphs into Memory](#).)

The database schema storing the graph must have the privileges listed in [Required Privileges for Database Users](#).

- [Two-Tier Deployments of Oracle Graph with Autonomous Database](#)

In two-tier deployments, the client graph application connects directly to the Autonomous Database.

- [Three-Tier Deployments of Oracle Graph with Autonomous Database](#)

In three-tier deployments, the client graph application connects to PGX in a middle tier, and PGX connects to the Autonomous Database.

1.10.1 Two-Tier Deployments of Oracle Graph with Autonomous Database

In two-tier deployments, the client graph application connects directly to the Autonomous Database.

1. Install Oracle Graph Client, as explained in [Installing the Java Client](#).
2. Establish a JDBC connection, as described in the [Oracle Autonomous Warehouse documentation](#).

You must download the wallet and unzip it to a secure location. You can then reference it when establishing the connection as shown in [Example 1-1](#).

3. Start the Java Shell as shown in the code:

```
/bin/opg4j --no_connect
```

4. Connect to your database as shown in [Example 1-1](#).

Note:

If you need to use the Graph Visualization Application, you must additionally install the Oracle Graph Server.

- See [Installing Oracle Graph Server](#) for more details.
- See [Deploying the Graph Visualization Application](#) for more details on deploying the Graph Visualization Application in Tomcat or Oracle WebLogic Server.

Example 1-1 Creating a Database Connection in a Two-Tier Graph Deployment with Autonomous Database

```
opg4j> var jdbcUrl = "jdbc:oracle:thin:@<tns_alias>?  
TNS_ADMIN=<wallet_location>" // jdbc url to the DB  
opg4j> var user = "<user>"  
opg4j> var pass = "<password>"  
opg4j> var conn = DriverManager.getConnection(jdbcUrl, user, pass) //  
connecting to the DB  
conn ==> oracle.jdbc.driver.T4CConnection@57e6cb01
```

In the preceding example:

- **<tns_alias>**: TNS alias used in `tnsnames.ora` file
- **<wallet_location>**: Path to the directory where the wallet is stored
- **<user>**: Name of the database user
- **<password>**: Password for the user

1.10.2 Three-Tier Deployments of Oracle Graph with Autonomous Database

In three-tier deployments, the client graph application connects to PGX in a middle tier, and PGX connects to the Autonomous Database.

The wallets downloaded from the Oracle Cloud Console are mainly *routing wallets*, meaning they are used to route the connection to the right database and to encrypt the connection. In most cases, they are not auto-login wallets, so they do not contain the password for the actual connection. The password usually needs to be provided separately to the wallet location.

The graph server does not support a wallet stored on the client file system or provided directly by remote users. The high level implications of this are:

- The server administrator provides the wallet and stores the wallet securely on the server's file system.
- Similar to Java EE connection pools, remote users will use that wallet when connecting. This means the server administrator trusts all remote users to use the wallet. As with any production deployments, the PGX server must be configured to enforce authentication and authorization to establish that trust.
- Remote users still need to provide a user name and password when sending a graph read request, just as with non-autonomous databases.
- You can only configure one wallet for each PGX server.

Having the same PGX server connecting to multiple Autonomous Databases is not supported. If you have that use case, start one PGX server for each Autonomous Database.

Pre-loaded graphs

To read a graph from Autonomous Database into PGX at server startup, follow the steps described in [Store the Database Password in a Keystore](#) to:

1. Create a Java Keystore containing the database password
2. Create a PGX graph configuration file describing the location and properties of the graph to be loaded
3. Update the `/opt/oracle/graph/pgx.conf` file to reference the graph configuration file

As root user, edit the service file at `/etc/systemd/system/pgx.service` and specify the environment variable under the `[Service]` directive:

```
Environment="JAVA_OPTS=-Doracle.net.tns_admin=/etc/oracle/graph/wallets"
```

Make sure that the directory (`/etc/oracle/graph/wallets` in the preceding code) is readable by the Oracle Graph user, which is the user that starts up the PGX server when using `systemd`.

In addition, edit the `ExecStart` command to specify the location of the keystore containing the password:

```
ExecStart=/bin/bash start-server --secret-store /etc/keystore.p12
```

 **Note:**

Please note that `/etc/keystore.p12` must not be password protected for this to work. Instead protect the file via file system permission that is only readable by `oraclegraph` user.

After the file is edited, reload the changes using:

```
systemctl daemon-reload
```

Finally start the server:

```
sudo systemctl start pgx
```

On-demand graph loading

To allow remote users of PGX to read from the Autonomous Database on demand, you can choose from two approaches:

- Provide the path to the wallet at server startup time via the `oracle.net.tns_admin` system property. Remote users have to provide the TNS address name, username and keystore alias (password) in their graph configuration files. The wallet is stored securely on the graph server's file system, and the server administrator trusts all remote users to use the wallet to connect to an Autonomous Database.

For example, the server administrator edits the service file at `/etc/systemd/system/pgx.service` and specifies the environment variable the under the `[Service]` directive:

```
Environment="JAVA_OPTS=-Doracle.net.tns_admin=/etc/oracle/graph/wallets"
```

and then start the server using

```
systemctl start pgx
```

The `/etc/oracle/graph/wallets/tnsnames.ora` file contains an address as follows:

```
sombrero_medium = (description= (retry_count=20) (retry_delay=3)
(address=(protocol=tcps) (port=1522) (host=adb.us-
ashburn-1.oraclecloud.com)
(connect_data=(service_name=l8lgholga0ujxsa_sombrero_medium.adwc.ora
clecloud.com)) (security=(ssl_server_cert_dn="CN=adwc.uscom-
```

```
east-1.oraclecloud.com,OU=Oracle BMCS US,O=Oracle Corporation,L=Redwood  
City,ST=California,C=US"))
```

Now remote users can read data into the server by sending a graph configuration file with the following connection properties:

```
{  
  ...  
  "jdbc_url": "jdbc:oracle:thin:@sombbrero_medium",  
  "username": "hr",  
  "keystore_alias": "database1",  
  ...  
}
```

Note that the keystore still lives on the client side and should contain the password for the `hr` user referenced in the config object, as explained in [Store the Database Password in a Keystore](#). A similar approach works for Tomcat or WebLogic Server deployments.

- Use Java EE connection pools in your web application server. Remote users only have to provide the name of the datasource in their graph configuration files. The wallet and the connection credentials are stored securely in the web application server's file system, and the server administrator trusts all remote users to use a connection from the pool to connect to an Autonomous Database.

You can find instructions how to set up such a data source at the following locations:

- WebLogic Server: [Configuring a WebLogic Data Source to use ATP](#)
- Tomcat: <https://www.oracle.com/technetwork/database/application-development/jdbc/documentation/atp-5073445.html#Tomcat>

If you gave the data source the name `adb_ds`, you can reference them by sending a graph configuration file with the following connection properties:

```
{  
  ...  
  "datasource_id": "adb_ds",  
  ...  
}
```

1.11 Migrating Property Graph Applications from Before Release 21c

If you are migrating from a previous version of Oracle Spatial and Graph to Release 21c, you may need to make some changes to existing property graph-related applications.

Also note that Oracle Graph Server and Client is required for property graph applications. This can be downloaded from [Oracle Software Delivery Cloud](#) or from [Oracle Downloads](#) page.

Security-Related Changes

The Property Graph feature contains a series of enhancements to further strengthen the security of the property graph component of product. The following enhancements may require manual changes to existing graph applications so that they continue to work properly.

- **Graph configuration files now require sensitive information such as passwords to be stored in Java Keystore files**
If you use graph configuration files you are required to use Java Keystore files to store sensitive information such as passwords. (See [Store the Database Password in a Keystore](#) for how to create and reference such a keystore.)
All existing graph configuration files with secrets in them must be migrated to the keystore-based approach.
- **In a three-tier deployment, access to the PGX server file system requires a directories allowlist**
By default, the PGX server does not allow remote access to the local file system. This can be explicitly allowed, though, in `/etc/oracle/graph/pgx.conf` by setting `allow_local_filesystem` to `true`. If you set `allow_local_filesystem` to `true`, you must also specify a list of directories that are allowed to be accessed, by setting `datasource_dir_whitelist`. For example:

```
"allow_local_filesystem": true,  
"datasource_dir_whitelist": ["/scratch/data1", "/scratch/data2"]
```

This will allow remote users to read and write data on the server's file-system from and into `/scratch/data1` and `/scratch/data2`.

- **In a three-tier deployment, reading from remote locations into PGX is no longer allowed by default**
Previously, PGX allowed graph data to be read from remote locations over FTP or HTTP. This is no longer allowed by default and requires explicit opt-in by the server administrator. To opt-in, specify the `allowed_remote_loading_locations` configuration option in `/etc/oracle/graph/pgx.conf`. For example:

```
allowed_remote_loading_locations: ["*"]
```

In addition:

- The `ftp` and `http` protocols are no longer supported for loading or storing data because they are unencrypted and thus insecure.
- Configuration files can no longer be loaded from remote locations, but must be loaded from the local file system.
- **Removed shell command line options**
The following command line options of the Groovy-based `opg` shell have been removed and will no longer work:
 - `--attach` - the shell no longer supports attaching to existing sessions via command line
 - `--password` - the shell will prompt now for the password

Also note that the Groovy-based shell has been deprecated, and you are encouraged to use the new JShell-based shell instead (see [Interactive Graph Shell CLIs](#)).

- **Changes to PGX APIs**

The following APIs no longer return graph configuration information:

- `ServerInstance#getGraphInfo()`
- `ServerInstance#getGraphInfos()`
- `ServerInstance#getServerState()`

The REST API now identifies collections, graphs, and properties by UUID instead of a name.

The namespaces for graphs and properties are session private by default now. This implies that some operations that would previously throw an exception due to a naming conflict could succeed now.

`PgxGraph#publish()` throws an exception now if a graph with the given name has been published before.

Migrating Data to a New Database Version

Oracle Graph Server and Client works with older database versions. (See [Database Compatibility and Restrictions](#) for information.) If as part of your upgrade you also upgraded your Oracle Database, you can migrate your existing graph data that was stored using the Oracle Property Graph format by invoking the following helper script in your database after the upgrade:

```
sqlplus> EXECUTE mdsys.opg.migrate_pg_to_current(graph_name=>'mygraph');
```

The preceding example migrates the property graph *mygraph* to the current database version.

Uninstalling Previous Versions of Property Graph Libraries

This is only necessary if you are using Oracle Database versions 12.2, 18c, or 19c.

Use of the Property Graph feature of Oracle Database now requires Oracle Graph Server and Client that is installed separately. After you have completed the Graph Server and Client installation, complete the preceding migration steps (if needed), and confirmed that everything is working well, it is recommended that you remove the binaries of **older** graph installations from your Oracle Database installation by performing the following un-install steps:

1. Make sure the Property Graph mid-tier components are not in use on the target database host. For example, ensure that there is no application running which uses any files under `$ORACLE_HOME/md/property_graph`. Examples of such an application are a running PGX server on the same host as the database or a client application that references the JAR files under `$ORACLE_HOME/md/property_graph/lib`.

It is **not** necessary to shut down the database to perform the uninstall. The Oracle database itself does not reference or use any files under `$ORACLE_HOME/md/property_graph`.

2. Remove the files under `$ORACLE_HOME/md/property_graph` on your database host. On Linux, you can copy the following helper script to your database host and run it with as the DBA operating system user: `/opt/oracle/graph/scripts/patch-opg-oracle-home.sh`

1.12 Upgrading From Graph Server and Client 20.4.x to 21.x

If you are upgrading from Graph Server and Client 20.4.x to 21.x version, you may need to create new roles in database and migrate authorization rules from `pgx.conf` file to the database. Also, starting from Graph Server and Client Release 21.1, TLS is enforced at the time of the RPM file installation.

One of the main enhancements of Graph Server and Client Release 21.1 is moving the graph access permissions from the `pgx.conf` file to the database. A new set of graph roles with default permissions are created automatically in the database, at the time of the PL/SQL packages installation. See [Table A-1](#) in the appendix for more details on the default mappings.

In order to comply with this feature you must perform the database actions explained in the following sections:

Creating additional roles in the database

The roles in the database with additional privileges are created when you install the 21.x PL/SQL packages in your database as part of the upgrade. If you are not able to install the PL/SQL packages, for example if you are using an Autonomous Database, see [User Authentication and Authorization](#) for more information on manually creating these roles in the database with the default set of privileges.

Migrating authorization rules

You must execute database `GRANTS` for user-added mappings contained in the `pgx.conf` file when upgrading to 21.x.

The following examples explain the various scenarios where migration of authorization rules may or may not apply.

Example 1-2 Migrating user-added mappings to database

To migrate the following user-added mappings in `pgx.conf` file:

```
...
"authorization": [{
  "pgx_role": "GRAPH_DEVELOPER",
  "pgx_permissions": [{
    "grant": "PGX_SESSION_ADD_PUBLISHED_GRAPH"
  }],
},
...

```

```
GRANT
```

```
GRANT PGX_SESSION_ADD_PUBLISHED_GRAPH TO GRAPH_DEVELOPER
```

Example 1-3 Migrating user-added file system authorization rules to database

To migrate the following user-added file system authorization rules in `pgx.conf` file:

```
...
"file_locations": [{
  "name": "my_hdfs_graph_data",
  "location": "hdfs:/data/graphs"
}],
"authorization": [{
  "pgx_role": "GRAPH_DEVELOPER",
  "pgx_permissions": [{
    "file_location": "my_hdfs_graph_data",
    "grant": "read"
  }],
}],
...
```

```
GRANT
```

```
CREATE OR REPLACE DIRECTORY my_hdfs_graph_data AS 'hdfs:/data/graphs'
GRANT READ ON DIRECTORY my_hdfs_graph_data TO GRAPH_DEVELOPER
```

Example 1-4 User-added graph authorization rules for preloaded graphs**Note:**

No migration required for user-added graph authorization rules for preloaded graphs.

You must not migrate user-added graph authorization rules for preloaded graphs (as shown in the following code) as these rules continue to be configured in `pgx.conf` file.

```
"preload_graphs": [{
  "path": "/data/my-graph.json",
  "name": "global_graph"
}],
"authorization": [{
  "pgx_role": "GRAPH_DEVELOPER",
  "pgx_permissions": [{
    "preloaded_graph": "global_graph",
    "grant": "read"
  }],
}],
...
```

Self-signed TLS certificate now generated upon RPM installation

In Graph Server and Client 21.x the RPM installation generates a self-signed certificate into `/etc/oracle/graph`, which the server uses to enable TLS by default.

According to security best practices, access to the certificate is restricted to the `oraclegraph` operating system user. The implication of this is that you no longer can start the graph server via the `/opt/oracle/graph/pgx/bin/start-server` script, even if your user is part of the `oraclegraph` group. Instead, manage the lifecycle of the graph server via `systemctl` commands. For example:

```
sudo systemctl start pgx
```

Another possible option is to change the ownership of the certificate as shown:

```
sudo chown <youruser> /etc/oracle/graph/server_key.pem
```

Turning off TLS is not recommended as it reduces the security of your connection. However, if you must do so, see [Disabling Transport Layer Security \(TLS\) in Graph Server](#) for more details.

1.13 Using the Graph Zeppelin Interpreter Client

Oracle Graph provides an interpreter client implementation for Apache Zeppelin. This tutorial topic explains how to perform simple operations using the graph Zeppelin interpreter client.

See [Installing the Graph Zeppelin Interpreter Client](#) for more details to install the graph interpreter into your local Zeppelin installation.

Using the Interpreter

If you named the graph interpreter `pgx`, you can send paragraphs to the graph server by starting the paragraphs with the `%pgx` directive, just as with any other interpreter.

The interpreter acts like a client that talks to a remote graph server. You cannot run a graph server instance embedded inside the Zeppelin interpreter. You must provide the graph server base URL and connection information as illustrated in the following example:

```
%pgx
import oracle.pgx.api.*
import groovy.json.*

baseUrl = '<base-url>'
username = '<username>'
password = '<password>'

conn = new URL("${baseUrl}/auth/token").openConnection()
conn.setRequestProperty('Content-Type', 'application/json')
token = conn.with {
    doOutput = true
    requestMethod = 'POST'
    outputStream.withWriter { writer ->
        writer << JsonOutput.toJson([username: username, password:
password])
    }
    return new JsonSlurper().parseText(content.text).access_token
```

```

}

instance = Pgx.getInstance(baseUrl, token)
session = instance.createSession("my-session")

```

The in-memory analyst Zeppelin interpreter evaluates paragraphs in the same way that the in-memory analyst shell does, and returns the output. Therefore, any valid in-memory analyst shell script will run in the in-memory analyst interpreter, as in the following example:

```

%pgx
g_brands = session.readGraphWithProperties("/opt/data/exommerce/
brand_cat.json")
g_brands.getNumVertices()
rank = analyst.pagerank(g_brands, 0.001, 0.85, 100)
rank.getTopKValues(10)

```

The following figure shows the results of that query after you click the icon to execute it.



ID	value
Cell Phones & Accessories	0.10107276500035282
Cases	0.060593137960391966
Basic Cases	0.058782080785810285
Accessories	0.05657872563693525

As you can see in the preceding figure, the Zeppelin interpreter automatically renders the values returned by `rank.getTopKValues(10)` as a Zeppelin table, to make it more convenient for you to browse results.

Besides the property values (`getTopKValues()`, `getBottomKValues()`, and `getValues()`), the following return types are automatically rendered as table also if they are returned from a paragraph:

- `PgsqlResultSet` - the object returned by the `queryPgsql("...")` method of the `PgxGraph` class.
- `MapIterable` - the object returned by the `entries()` method of the `PgxMap` class

All other return types and errors are returned as normal strings, just as the in-memory analyst shell does.

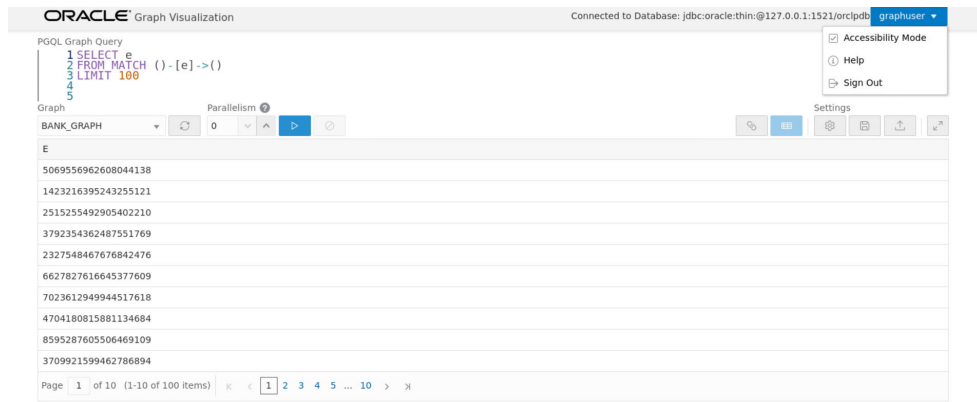
For more information about Zeppelin, see the [official Zeppelin documentation](#).

1.14 About Oracle Graph Server and Client Accessibility

This section provides information on the accessibility features for Oracle Graph Server and Client.

- For information on addressing accessibility for the Java Client and the Python Client which are installed on Oracle Linux, see [Working With Accessibility Features in Oracle Linux 7](#).
- For information on addressing accessibility for the Graph Visualization Application, which is based on Oracle JET, see [About Oracle JET and Accessibility](#).
- You can enable the accessibility mode for the Graph Visualization user interface by clicking the user menu on the top-right and selecting the **Accessibility Mode** checkbox. Once enabled, the query output is always displayed in a tabular layout as shown:

Figure 1-7 Enabling Accessibility Mode in the Graph Visualization Application



2

Quick Starts for Using Oracle Property Graph

This chapter contains quick start tutorials and other resources to help you get started on working with Oracle property graphs.

- [Quick Start: Interactively Analyze Graph Data](#)
This tutorial shows how you can quickly get started using property graph data and learn to execute PGQL queries and run graph algorithms on the data and display results.
- [QuickStart: Run Graph Analytics Using the Python Shell](#)
This tutorial shows how you can get started using property graph data using the Python shell.
- [Quick Start: Using the Python Client as a Module](#)
This section describes how to use the Python client as a module in Python applications.
- [Oracle LiveLabs Workshops for Graphs](#)
You can also explore Oracle Property Graph features using the graph workshops in Oracle LiveLabs.

2.1 Quick Start: Interactively Analyze Graph Data

This tutorial shows how you can quickly get started using property graph data and learn to execute PGQL queries and run graph algorithms on the data and display results.

The tutorials in this section are:

- [Quick Start: Create and Query a Graph in the Database, Load into In-Memory Graph Server \(PGX\) for Analytics](#)
This tutorial shows how you can get started using property graph data when you create a graph and persist it in the database. The graph can be queried in the database. This tutorial uses the JShell client.
- [Quick Start: Create, Query, and Analyze a Graph in In-Memory Graph Server \(PGX\)](#)
This tutorial shows how you can quickly get started using property graph data when using the in-memory graph server (PGX).
- [Quick Start: Executing PGQL Queries in SQLcl](#)
This tutorial provides you resources to get started on executing PGQL Queries in SQLcl.

2.1.1 Quick Start: Create and Query a Graph in the Database, Load into In-Memory Graph Server (PGX) for Analytics

This tutorial shows how you can get started using property graph data when you create a graph and persist it in the database. The graph can be queried in the database. This tutorial uses the JShell client.

See [Create and Query a Graph in the Database](#) for more information on creating and storing graphs in database.

- Convert existing relational data into a graph in the database.

- Query this graph using PGQL.

In [Load the Graph into Memory and Run Graph Analytics](#), you will run graph algorithms after loading the graph into the in-memory graph server (PGX).

- Load the graph into the in-memory graph server (PGX), run graph algorithms on this graph, and visualize results.

Prerequisites for the following quickstart are:

- An installation of Oracle Graph server (this is PGX, the in-memory graph server). See [Oracle Graph Server and Client Installation](#) for information to download Oracle Graph Server and Client.

- An installation of Oracle Graph client

- Java 11

- The in-memory graph server can work with Java 8 or Java 11.
- The JShell client used in this example requires Java 11.

For Java downloads, see <https://www.oracle.com/technetwork/java/javase/overview/index.html>.

- Connection details for your Oracle Database. See [Database Compatibility and Restrictions](#) to identify any limitations. The Property Graph feature is supported for Oracle Database versions 12.2 and later.
- Basic knowledge about how to run commands on Oracle Database (for example, using `SQL*Plus` or `SQL Developer`).

Set up the example data

This example uses the HR (human resources) sample dataset.

- For instructions how to import that data into a user managed database, see: <https://github.com/oracle/db-sample-schemas>
- If you are using Autonomous Database, see: <https://www.thatjeffsmith.com/archive/2019/07/creating-hr-in-oracle-autonomous-database-w-sql-developer-web/>

Note that the database schema storing the graph must have the privileges listed in [Required Privileges for Database Users](#).

- [Create and Query a Graph in the Database](#)
In this section, you will use the Oracle Graph client to create a graph from relational tables and store it in the property graph schema in the database.
- [Load the Graph into Memory and Run Graph Analytics](#)

2.1.1.1 Create and Query a Graph in the Database

In this section, you will use the Oracle Graph client to create a graph from relational tables and store it in the property graph schema in the database.

Major tasks for this tutorial:

- [Start the shell](#)
- [Open a JDBC database connection](#)
- [Create a PGQL connection](#)

- [Write and execute the graph creation statement](#)
- [Run a few PGQL queries](#)

Start the shell

On the system where Oracle Graph client is installed, start the shell by as follows:

```
cd <client-install-dir>
./bin/opg4j --noconnect
```

The `--noconnect` option indicates that you are not connecting to the in-memory graph server (PGX). You will only be connecting to the database in this example.

Note that `JAVA_HOME` should be set to Java 11 before you start the shell. For example:

```
export JAVA_HOME=/usr/lib/jvm/java-11-oracle
```

See [Interactive Graph Shell CLIs](#) for details about the shell.

Open a JDBC database connection

Inside the shell prompt, use the standard JDBC Java API to obtain a database connection object. For example:

```
opg4j> var jdbcUrl = "<jdbc-url>" // for example:
jdbc:oracle:thin:@myhost:1521/myervice
opg4j> var user = "<db-user>" // for example: hr
opg4j> var pass = "<db-pass>"
opg4j> var conn = DriverManager.getConnection(jdbcUrl, user, pass)
conn ==> oracle.jdbc.driver.T4CConnection@57e6cb01
```

Connecting to an Autonomous Database works the same way: provide a JDBC URL that points to the local wallet. See [Using Oracle Graph with the Autonomous Database](#) for an example.

Create a PGQL connection

Convert the JDBC connection into a PGQL connection object. For example:

```
opg4j> conn.setAutoCommit(false)
opg4j> var pgql = PgqlConnection.getConnection(conn)
pgql ==> oracle.pg.rdbms.pgql.PgqlConnection@6fb3d3bb
```

Write and execute the graph creation statement

Using a text editor, write a `CREATE PROPERTY GRAPH` statement that describes how the HR sample data should be converted into a graph. Save this file as `create.pgql` at a location of your choice. For example:

```
CREATE PROPERTY GRAPH hr
  VERTEX TABLES (
    employees LABEL employee
    PROPERTIES ARE ALL COLUMNS EXCEPT ( job_id, manager_id,
```

```

department_id ),
  departments LABEL department
    PROPERTIES ( department_id, department_name ),
  jobs LABEL job
    PROPERTIES ARE ALL COLUMNS,
  job_history
    PROPERTIES ( start_date, end_date ),
  locations LABEL location
    PROPERTIES ARE ALL COLUMNS EXCEPT ( country_id ),
  countries LABEL country
    PROPERTIES ARE ALL COLUMNS EXCEPT ( region_id ),
  regions LABEL region
)
EDGE TABLES (
  employees AS works_for
    SOURCE employees
    DESTINATION KEY ( manager_id ) REFERENCES employees
    NO PROPERTIES,
  employees AS works_at
    SOURCE employees
    DESTINATION departments
    NO PROPERTIES,
  employees AS works_as
    SOURCE employees
    DESTINATION jobs
    NO PROPERTIES,
  departments AS managed_by
    SOURCE departments
    DESTINATION employees
    NO PROPERTIES,
  job_history AS for_employee
    SOURCE job_history
    DESTINATION employees
    LABEL for
    NO PROPERTIES,
  job_history AS for_department
    SOURCE job_history
    DESTINATION departments
    LABEL for
    NO PROPERTIES,
  job_history AS for_job
    SOURCE job_history
    DESTINATION jobs
    LABEL for
    NO PROPERTIES,
  departments AS department_located_in
    SOURCE departments
    DESTINATION locations
    LABEL located_in
    NO PROPERTIES,
  locations AS location_located_in
    SOURCE locations
    DESTINATION countries
    LABEL located_in
    NO PROPERTIES,

```

```

        countries AS country_located_in
        SOURCE countries
        DESTINATION regions
        LABEL located_in
        NO PROPERTIES
    )

```

Then, back in your graph shell, execute the `CREATE PROPERTY GRAPH` statement by sending it to your PGQL connection. Replace `<path>` with the path to the directory containing the `create.pgql` file:

```

opg4j> pgql.prepareStatement(Files.readString(Paths.get("<path>/
create.pgql"))).execute()
$16 ==> false

```

Run a few PGQL queries

Now that you have a graph named `hr`, you can use PGQL to run a few queries against it directly on the database. For example:

```

// define a little helper function that executes the query, prints the
// results and properly closes the statement
opg4j> Consumer<String> query = q -> { try(var s = pgql.prepareStatement(q))
{ s.execute(); s.getResultSet().print(); } catch(Exception e) { throw new
RuntimeException(e); } }
query ==> $Lambda$605/0x0000000100ae6440@6c9e7af2

// print the number of vertices in the graph
opg4j> query.accept("select count(v) from hr match (v)")
+-----+
| count(v) |
+-----+
| 215      |
+-----+

// print the number of edges in the graph
opg4j> query.accept("select count(e) from hr match ()-[e]->()")
+-----+
| count(e) |
+-----+
| 433      |
+-----+

// find the highest earning managers
opg4j> query.accept("select distinct m.FIRST_NAME, m.LAST_NAME, m.SALARY
from hr match (v:EMPLOYEE)-[:WORKS_FOR]->(m:EMPLOYEE) order by m.SALARY
desc")
+-----+
| m.FIRST_NAME | m.LAST_NAME | m.SALARY |
+-----+
| Steven       | King        | 24000.0  |
| Lex          | De Haan     | 17000.0  |
| Neena        | Kochhar     | 17000.0  |
| John         | Russell    | 14000.0  |

```

```
| Karen          | Partners  | 13500.0 |
| Michael       | Hartstein | 13000.0 |
| Alberto       | Errazuriz | 12000.0 |
| Shelley       | Higgins   | 12000.0 |
| Nancy         | Greenberg | 12000.0 |
| Den           | Raphaely  | 11000.0 |
| Gerald        | Cambrault | 11000.0 |
| Eleni         | Zlotkey   | 10500.0 |
| Alexander     | Hunold    | 9000.0  |
| Adam          | Fripp     | 8200.0  |
| Matthew       | Weiss     | 8000.0  |
| Payam         | Kaufling  | 7900.0  |
| Shanta        | Vollman   | 6500.0  |
| Kevin         | Mourgos   | 5800.0  |
+-----+
```

```
// find the average salary of accountants in the Americas
opg4j> query.accept("select avg(e.SALARY) from hr match (e:EMPLOYEE) -
[h:WORKS_AT]-> (d:DEPARTMENT) -[:LOCATED_IN]-> (:LOCATION) -
[:LOCATED_IN]-> (:COUNTRY) -[:LOCATED_IN]-> (r:REGION) where
r.REGION_NAME = 'Americas' and d.DEPARTMENT_NAME = 'Accounting'")
+-----+
| avg(e.SALARY) |
+-----+
| 14500.0       |
+-----+
```

2.1.1.2 Load the Graph into Memory and Run Graph Analytics

Major tasks for this tutorial:

- [Load the graph from the property graph schema into memory](#)
- [Execute algorithms and query the algorithm results](#)
- [Share the Graph with Other Sessions](#)

Load the graph from the property graph schema into memory

In this section of the quickstart, you will load the graph stored in the Property Graphs schema in the database into the in-memory graph server (PGX). This will enable you to run a variety of different built-in algorithms on the graph and will also improve query performance for larger graphs.

First, start the JShell client and connect to the in-memory graph server (PGX):

```
./bin/opg4j --base_url https://<graph server host>:7007 --username
<graphuser>
```

<graphuser> is the database user you will use to for the PGX server authentication. You will be prompted for the database password.

 **Note:**

For demo purposes only, if you have set `enable_tls` to `false` in the `/etc/oracle/graph/server.conf` file you can use an `http` instead of `https` connection.

```
./bin/opg4j --base_url http://<graph server host>:7007 --username <graphuser>
```

This starts the shell and makes a connection to the graph server.

 **Note:**

Always use low-privilege read-only database user accounts for PGX, as explained in [Security Best Practices with Graph Data](#).

Next load the graph into memory in this server.

To load the graph into memory, create a PGX graph config object, using the PGX graph config builder API to do this directly in the shell.

The following example creates a PGX graph config object. It lists the properties to load into memory so that you can exclude other properties, thus reducing memory consumption.

```
Supplier<GraphConfig> pgxConfig = () -> { return
GraphConfigBuilder.forPropertyGraphRdbms()
.setName("hr")
.addVertexProperty("COUNTRY_NAME", PropertyType.STRING)
.addVertexProperty("DEPARTMENT_NAME", PropertyType.STRING)
.addVertexProperty("FIRST_NAME", PropertyType.STRING)
.addVertexProperty("LAST_NAME", PropertyType.STRING)
.addVertexProperty("EMAIL", PropertyType.STRING)
.addVertexProperty("PHONE_NUMBER", PropertyType.STRING)
.addVertexProperty("SALARY", PropertyType.DOUBLE)
.addVertexProperty("MIN_SALARY", PropertyType.DOUBLE)
.addVertexProperty("MAX_SALARY", PropertyType.DOUBLE)
.addVertexProperty("STREET_ADDRESS", PropertyType.STRING)
.addVertexProperty("POSTAL_CODE", PropertyType.STRING)
.addVertexProperty("CITY", PropertyType.STRING)
.addVertexProperty("STATE_PROVINCE", PropertyType.STRING)
.addVertexProperty("REGION_NAME", PropertyType.STRING)
.setPartitionWhileLoading(PartitionWhileLoading.BY_LABEL)
.setLoadVertexLabels(true)
.setLoadEdgeLabel(true)
.build(); }
```

Now that you have a graph config object, use the following API to read the graph into PGX:

```
opg4j> var graph = session.readGraphWithProperties(pgxConfig.get())
graph ==> PgxGraph[name=hr,N=215,E=433,created=1586996113457]
```

The session object is created for you automatically.

Execute algorithms and query the algorithm results

Now that you have the graph in memory, you can run any built-in algorithm using a single API invocation. For example, for `pagerank`:

```
opg4j> analyst.pagerank(graph)
$31==> VertexProperty[name=pagerank,type=double,graph=hr]
```

As you can see from the preceding outputs, each algorithm created a new vertex property on the graph holding the output of the algorithm. To print the most important people in the graph (according to `pagerank`), you can run the following query:

```
opg4j> session.queryPgql("select m.FIRST_NAME, m.LAST_NAME, m.pagerank
from hr match (m:EMPLOYEE) order by m.pagerank desc limit
10").print().close()
+-----+
| m.FIRST_NAME | m.LAST_NAME | m.pagerank |
+-----+
| Adam | Fripp | 0.002959240305566317 |
| John | Russell | 0.0028810951120575284 |
| Michael | Hartstein | 0.002181365227465801 |
| Alexander | Hunold | 0.002082616009054747 |
| Den | Raphaely | 0.0020378615199327507 |
| Shelley | Higgins | 0.002028946863425767 |
| Nancy | Greenberg | 0.0017419394483596667 |
| Steven | King | 0.0016622985848193119 |
| Neena | Kochhar | 0.0015252785582170803 |
| Jennifer | Whalen | 0.0014263044976976823 |
+-----+
```

Share the Graph with Other Sessions

After you load the graph into the in-memory graph server, you can use the `publish()` API to make the graph visible to other sessions, such as the graph visualization session. For example:

```
opg4j> graph.publish(VertexProperty.ALL, EdgeProperty.ALL)
```

The published graph will include any new properties you add to the graph by calling functions, such as `pagerank`.

You can use the [Graph Visualization Application](#) by navigating to `<my-server-name>:7007/ui/` in your browser.

You can connect to a particular client session by providing the session ID when you log into the Graph Visualization Application. You will then be able to visualize all graphs in the session, even if they have not been published.

```
opg4j> session
session ==> PgxSession[ID=5adf83ab-31b1-4a0e-8c08-
d6a95ba63ee0,source=pgxShell]
```

The session id is 5adf83ab-31b1-4a0e-8c08-d6a95ba63ee0.



Note:

You must create a server certificate to connect to the in-memory graph server (PGX) from the Graph Visualization Application. See [Setting Up Transport Layer Security](#) for more details.

2.1.2 Quick Start: Create, Query, and Analyze a Graph in In-Memory Graph Server (PGX)

This tutorial shows how you can quickly get started using property graph data when using the in-memory graph server (PGX).

This is for use cases where the graph is available as long as the in-memory graph server (PGX) session is active. The graph is not persisted in the database.

- Create a graph in the in-memory graph server (PGX), directly from existing relational data
- Query this graph using PGQL in the in-memory graph server (PGX)
- Run graph algorithms in the in-memory graph server (PGX) on this graph and display results

Prerequisites for the following quickstart are:

- An installation of Oracle Graph server (this is PGX, the in-memory graph server).
See [Installing Oracle Graph Server](#) for information to download Oracle Graph Server.
- An installation of Oracle Graph client.
See [Installing the Java Client](#) for information to download Oracle Graph Client.

You will authenticate yourself as the database user to the in-memory graph server, and these database credentials are used to access the database tables and create a graph.

- Java 11
 - The in-memory graph server can work with Java 8 or Java 11.
 - The JShell client used in this example requires Java 11.

For Java downloads, see <https://www.oracle.com/technetwork/java/javase/overview/index.html>.

Major tasks for this tutorial:

- [Set up the example data](#)
- [Start the shell](#)
- [Write and execute the graph creation statement](#)
- [Run a few PGQL queries](#)
- [Execute algorithms and query the algorithm results](#)
- [Share the Graph with Other Sessions](#)

Set up the example data

This example uses the HR (human resources) sample dataset.

- For instructions how to import that data into a user managed database, see: <https://github.com/oracle/db-sample-schemas>
- If you are using Autonomous Database, see: <https://www.thatjeffsmith.com/archive/2019/07/creating-hr-in-oracle-autonomous-database-w-sql-developer-web/>

Note that the database schema storing the graph must have the privileges listed in [Required Privileges for Database Users](#).

Start the shell

On the system where Oracle Graph Client is installed, start the shell as follows. This is an example of starting a shell in remote mode and connecting to the in-memory graph server (PGX):

```
./bin/opg4j --base_url https://<graph server host>:7007 --username  
<graphuser>
```

<graphuser> is the database user you will use to for the PGX server authentication. You will be prompted for the database password.



Note:

For demo purposes only, if you have set `enable_tls` to `false` in the `/etc/oracle/graph/server.conf` file you can use an `http` instead of `https` connection.

```
./bin/opg4j --base_url http://<graph server host>:7007 --username  
<graphuser>
```

This starts the shell and makes a connection to the graph server.

Note that, `JAVA_HOME` should be set to Java 11 before you start the shell. For example:

```
export JAVA_HOME=/usr/lib/jvm/java-11-oracle
```

See [Interactive Graph Shell CLIs](#) for details about the shell.

Write and execute the graph creation statement

Create a graph with employees, departments, and “employee works at department”, by executing a `CREATE PROPERTY GRAPH` statement. The following statement creates a graph in the in-memory graph server (PGX):

```
opg4j> String statement =  
    "CREATE PROPERTY GRAPH hr_simplified "  
    + " VERTEX TABLES ( "  
    + "   hr.employees LABEL employee "
```

```

+ "      PROPERTIES ARE ALL COLUMNS EXCEPT ( job_id, manager_id,
department_id ), "
+ "      hr.departments LABEL department "
+ "      PROPERTIES ( department_id, department_name ) "
+ " ) "
+ " EDGE TABLES ( "
+ "      hr.employees AS works_at "
+ "      SOURCE KEY ( employee_id ) REFERENCES employees "
+ "      DESTINATION departments "
+ "      PROPERTIES ( employee_id ) "
+ " )";
opg-jshell> session.executePgql(statement);

```

To get a handle to the graph, execute:

```
opg4j> PgxGraph g = session.getGraph("HR_SIMPLIFIED");
```

Run a few PGQL queries

You can use this handle to run PGQL queries on this graph. For example, to find the department that “Nandita Sarchand” works for, execute:

```

opg4j> String query =
    "SELECT dep.department_name "
+ "FROM MATCH (emp:Employee) -[:works_at]-> (dep:Department) "
+ "WHERE emp.first_name = 'Nandita' AND emp.last_name = 'Sarchand' "
+ "ORDER BY 1";
opg4j> PgqlResultSet resultSet = g.queryPgql(query);
opg4j> resultSet.print();
+-----+
| department_name |
+-----+
| Shipping        |
+-----+

```

To get an overview of the types of vertices and their frequencies, execute:

```

opg4j> String query =
    "SELECT label(n), COUNT(*) "
+ "FROM MATCH (n) "
+ "GROUP BY label(n) "
+ "ORDER BY COUNT(*) DESC";
opg4j> PgqlResultSet resultSet = g.queryPgql(query);
opg4j> resultSet.print();
+-----+
| label(n)  | COUNT(*) |
+-----+
| EMPLOYEE  | 107      |
| DEPARTMENT | 27       |
+-----+

```

To get an overview of the types of edges and their frequencies, execute:

```
opg4j> String query =
    "SELECT label(n) AS srcLbl, label(e) AS edgeLbl, label(m) AS
dstLbl, COUNT(*) "
    + "FROM MATCH (n) -[e]-> (m) "
    + "GROUP BY srcLbl, edgeLbl, dstLbl "
    + "ORDER BY COUNT(*) DESC";
opg4j> PgqlResultSet resultSet = g.queryPgql(query);
opg4j> resultSet.print();
```

```
+-----+
| srcLbl | edgeLbl | dstLbl | COUNT(*) |
+-----+
| EMPLOYEE | WORKS_AT | DEPARTMENT | 106 |
+-----+
```

Execute algorithms and query the algorithm results

Now that you have the graph in memory, you can run each built-in algorithms using a single API invocation. For example, for pagerank:

```
opg4j> analyst.pagerank(g)
$31==> VertexProperty[name=pagerank,type=double,graph=hr]
```

As you can see from the preceding outputs, each algorithm created a new vertex property on the graph holding the output of the algorithm. To print the most important people in the graph (according to pagerank), you can run the following query:

```
opg4j> session.queryPgql("select m.FIRST_NAME, m.LAST_NAME, m.pagerank
from HR_SIMPLIFIED match (m:EMPLOYEE) where m.FIRST_NAME = 'Nandita'
").print().close()
```

```
+-----+
| m.FIRST_NAME | m.LAST_NAME | m.pagerank |
+-----+
| Nandita | Sarchand | 0.001119402985074627 |
+-----+
```

In the following example, we order departments by their pagerank value. Departments with higher pagerank values have more employees.

```
opg4j> session.queryPgql("select m.DEPARTMENT_NAME, m.pagerank from
HR_SIMPLIFIED match (m:DEPARTMENT) order by m.pagerank
").print().close();
```

```
+-----+
| m.DEPARTMENT_NAME | m.pagerank |
+-----+
| Manufacturing | 0.001119402985074627 |
| Construction | 0.001119402985074627 |
| Contracting | 0.001119402985074627 |
| Operations | 0.001119402985074627 |
```

```
| IT Support | 0.001119402985074627 |
| NOC | 0.001119402985074627 |
| IT Helpdesk | 0.001119402985074627 |
| Government Sales | 0.001119402985074627 |
| Retail Sales | 0.001119402985074627 |
| Recruiting | 0.001119402985074627 |
| Payroll | 0.001119402985074627 |
| Treasury | 0.001119402985074627 |
| Corporate Tax | 0.001119402985074627 |
| Control And Credit | 0.001119402985074627 |
| Shareholder Services | 0.001119402985074627 |
| Benefits | 0.001119402985074627 |
| Human Resources | 0.0020708955223880596 |
| Administration | 0.0020708955223880596 |
| Public Relations | 0.0020708955223880596 |
| Marketing | 0.003022388059701493 |
| Accounting | 0.003022388059701493 |
| Executive | 0.003973880597014925 |
| IT | 0.005876865671641792 |
| Purchasing | 0.006828358208955224 |
| Finance | 0.006828358208955224 |
| Sales | 0.03347014925373134 |
| Shipping | 0.043936567164179076 |
+-----+
```

Share the Graph with Other Sessions

After you load the graph into the server, you can use the `publish()` API to make the graph visible to other sessions, such as the graph visualization session. For example:

```
opg4j> graph.publish(VertexProperty.ALL, EdgeProperty.ALL)
```

The published graph will include any new properties you add to the graph by calling functions, such as `pagerank`.

Ensure that the logged-in user has the privilege to publish graphs. You can do this by adding the privilege `PGX_SESSION_ADD_PUBLISHED_GRAPH` to the `GRAPH_DEVELOPER` role as explained in [Adding Permissions to Publish the Graph](#). We had given the `GRAPH_DEVELOPER` role to the database user in [Installing PL/SQL Packages in Oracle Database](#).

You can use the Graph Visualization Application by navigating to `<my-server-name>:7007/ui/` in your browser.

You can connect to a particular client session by providing the session ID when you log into the Graph Visualization Application. You will then be able to visualize all graphs in the session, even if they have not been published.

```
opg4j> session
session ==> PgxCSession[ID=5adf83ab-31b1-4a0e-8c08-
d6a95ba63ee0,source=pgxShell]
```

The session id is `5adf83ab-31b1-4a0e-8c08-d6a95ba63ee0`.

**Note:**

You must create a server certificate to connect to the in-memory graph server (PGX) from the Graph Visualization Application. See [Setting Up Transport Layer Security](#) for more details.

2.1.3 Quick Start: Executing PGQL Queries in SQLcl

This tutorial provides you resources to get started on executing PGQL Queries in SQLcl.

You can execute PGQL queries in SQLcl with a plugin that is available with Oracle Graph Server and Client.

See [Execute PGQL Queries in SQLcl](#) for more details.

You can also refer to PGQL Plug-in for SQLcl [PGQL Plug-in for SQLcl](#) section in the SQLcl documentation.

2.2 QuickStart: Run Graph Analytics Using the Python Shell

This tutorial shows how you can get started using property graph data using the Python shell.

As a prerequisite for this quick start, you must ensure that you have completed the following installations:

- [Installing Oracle Graph Server](#)
- [Installing the Python Client](#)

1. Start the Python shell as shown:

```
./bin/opg4py --base_url https://localhost:7007
```

You are prompted to enter your username and password.

2. Verify that the Python client is connected to a remote graph server (PGX) instance as shown:

```
Oracle Graph Server Shell 21.4.0
>>> instance
ServerInstance(embedded: False, base_url: https://localhost:7007,
version: <oracle.pgx.common.VersionInfo at 0x7fb71a1b2f68
jclass=oracle/pgx/common/VersionInfo jself=<LocalRef obj=0xadd938
at 0x7fb71a1808f0>>>)
```

3. Create the graph using the graph builder Python API.

```
>>> graph = session.create_graph_builder().add_edge(1,
2).add_edge(2, 3).build("my_graph")
```

- Execute any built-in algorithm on the graph. For example:

```
>>> analyst.pagerank(graph)
VertexProperty(name: pagerank, type: double, graph: my_graph)
```

- Execute any PGQL queries and print the PGQL result set as shown:

```
>>> rs = session.query_pgql("select id(x), x.pagerank from match (x) on
my_graph")
>>> rs.print()
+-----+
| id(x) | pagerank          |
+-----+
| 1     | 0.0500000000000001 |
| 2     | 0.0925000000000003 |
| 3     | 0.1286250000000002 |
+-----+
```

Converting PGQL result set into pandas dataframe

Additionally, you can also convert the PGQL result set to a `pandas.DataFrame` object using the `to_pandas()` method. This makes it easier to perform various data filtering operations on the result set and it can also be used in Lambda functions. For example,

```
example_query = (
    "SELECT n.name as name, n.age as age "
    "WHERE (n) "
)
result_set = sample_graph.query_pgql(example_query)
result_df = result_set.to_pandas()

result_df['age_bin'] = result_df['age'].apply(lambda x: int(x)/20) #
create age bins based on age ranges
```



See Also:

[Python API Reference](#) to view the complete set of available Python APIs

2.3 Quick Start: Using the Python Client as a Module

This section describes how to use the Python client as a module in Python applications.

Embedded Server

You can use the python client as a module as illustrated in the following example.

 **Note:**

For this mode, the Python client and the Graph Server RPM package must be installed on the same machine.

```
import os
os.environ["PGX_CLASSPATH"] = "/opt/oracle/graph/lib/*"

import pypgx

session = pypgx.get_session()
graph = session.create_graph_builder().add_edge(1, 2).add_edge(2,
3).build("my_graph")
analyst = session.create_analyst()
analyst.pagerank(graph)
rs = session.query_pgql("select id(x), x.pagerank from match (x) on
my_graph")
rs.print()
```

To execute, save the above program into a file named `program.py` and run the following command.

```
python3 program.py
```

You will see the following output:

```
+-----+
| id(x) | pagerank          |
+-----+
| 1      | 0.0500000000000001 |
| 2      | 0.0925000000000003 |
| 3      | 0.1286250000000002 |
+-----+
```

See [Converting PGQL result set into pandas dataframe](#) for more details on converting a PGQL result set into pandas dataframe.

 **Note:**

To view the complete set of available Python APIs, see [Pypgx API](#).

Remote Server

For this mode, all you need is the Python client to be installed. In your Python program, you must authenticate with the remote server before you can create a session as illustrated in the following example.

**Note:**

Replace the `base_url`, `username`, and `password` with values to match your environment details.

```
import pypgx as pgx
import pypgx.pg.rdbms.graph_server as graph_server

base_url = "https://localhost:7007"
username = "scott"
password = "<password-for-scott>"

instance = graph_server.get_instance(base_url, username, password)
session = instance.create_session("python_pgx_client")
print(session)
```

To execute, save the above program into a file named `program.py` and run the following command:

```
python3 program.py
```

After successful login, you'll see the following message indicating a PGX session was created:

```
PgxSession(id: 0bdd4828-c3cc-4cef-92c8-0fcd105416f0, name: python_pgx_client)
```

**Note:**

To view the complete set of available Python APIs, see [Pypgx API](#).

2.4 Oracle LiveLabs Workshops for Graphs

You can also explore Oracle Property Graph features using the graph workshops in Oracle LiveLabs.

See the Oracle LiveLabs Workshop for a complete example on querying, analyzing and visualizing graphs using data stored in a free tier Autonomous Database instance. You will provision a new free tier Autonomous Database instance, load data into it, create a graph, and then query, analyze and visualize the graph.

3

Property Graph Views on Oracle Database Tables

You can create property graph views over data stored in Oracle Database. You can perform various graph analytics operations using PGQL on these views.

The `CREATE PROPERTY GRAPH` statement in PGQL can be used to create a view-like object that contains metadata about the graph. This graph can be queried using PGQL.

The property graph views are created directly over data that exists in the relational database tables. Since the graph is stored in the database tables it has a schema. This is unlike the graphs created with a flexible schema, where the data is copied from the source tables to property graph schema tables as described in [Property Graph Schema Objects for Oracle Database](#).

One of the main benefits of property graph views, is that all updates to the database tables are immediately reflected in the graph.

Metadata Tables for PG Views

Each time a `CREATE PROPERTY GRAPH` statement is executed, metadata tables are created in the user's own schema.

The following table describes the set of metadata tables that are created for each graph on executing `CREATE PROPERTY GRAPH` statement.

All columns shown underlined in the [Table 3-1](#) are part of the primary key of the table. Also all columns have a `NOT NULL` constraint.

Table 3-1 Metadata Tables for PG Views

Table Name	Description
<u>graphName_ELEM_TAB</u> <u>LE\$</u>	Metadata for graph element (vertex/edge) tables (one row per element table): <ul style="list-style-type: none"><u>ET_NAME</u>: the name of the element table (the "alias")<u>ET_TYPE</u>: either "VERTEX" or "EDGE"<u>SCHEMA_NAME</u>: the name of the schema of the underlying table<u>TABLE_NAME</u>: the name of underlying table
<u>graphName_LABEL\$</u>	Metadata on labels of element tables (one row per label; one label per element table): <ul style="list-style-type: none"><u>LABEL_NAME</u>: the name of the label<u>ET_NAME</u>: the name of the element table (the "alias")<u>ET_TYPE</u>: either "VERTEX" or "EDGE"
<u>graphName_PROPERTY</u> <u>\$</u>	Metadata describing the columns that are exposed through a label (one row per property) <ul style="list-style-type: none"><u>PROPERTY_NAME</u>: the name of the property<u>ET_NAME</u>: the name of the element table (the "alias")<u>ET_TYPE</u>: either "VERTEX" or "EDGE"<u>LABEL_NAME</u>: the name of the label that this property belongs to<u>COLUMN_NAME</u>: the name of the column (initially, only the case where property names equal column names is allowed)

Table 3-1 (Cont.) Metadata Tables for PG Views

Table Name	Description
graphName_ KEY \$	<p>Metadata describing a vertex/edge key (one row per column in the key)</p> <ul style="list-style-type: none"> <u>COLUMN_NAME</u>: the name of the column in the key <u>COLUMN_NUMBER</u>: the number of the column in the key For example, in KEY (a, b, c), "a" has number 1, "b" has number 2 and "c" has number 3. <u>KEY_TYPE</u>: either "VERTEX" or "EDGE" <u>ET_NAME</u>: the name of the element table (the "alias")
graphName_ SRC_DST _ KEY \$	<p>Metadata describing the edge source/destination keys (one row per column of a key):</p> <ul style="list-style-type: none"> <u>ET_NAME</u>: the name of the element table (the "alias"), which is always an edge table <u>VT_NAME</u>: the name of the vertex table <u>KEY_TYPE</u>: either "EDGE_SOURCE" or "EDGE_DESTINATION" <u>ET_COLUMN_NAME</u>: the name of the key column <u>ET_COLUMN_NUMBER</u>: the number of the column in the key. For example, in KEY (a, b, c), "a" has number 1, "b" has number 2 and "c" has number 3.

 **Note:**

Currently, support is only for **SOURCE KEY (...) REFERENCES T1**. So only the edge source/destination key is stored.

Example 3-1 To create a Property Graph View

Consider the following CREATE PROPERTY GRAPH statement:

```
CREATE PROPERTY GRAPH student_network
  VERTEX TABLES(
    person
      KEY ( id )
      LABEL student
      PROPERTIES( name ),
    university
      KEY ( id )
      PROPERTIES( name )
  )
  EDGE TABLES(
    knows
      key (person1, person2)
      SOURCE KEY ( person1 ) REFERENCES person
      DESTINATION KEY ( person2 ) REFERENCES person
      NO PROPERTIES,
    person AS studentOf
      key (id, university)
      SOURCE KEY ( id ) REFERENCES person
```

```

        DESTINATION KEY ( university ) REFERENCES university
        NO PROPERTIES
    )
    OPTIONS (PG_VIEW)

```

The `OPTIONS` clause allows the creation of a property graph view instead of the creation of property graph schema graph. You must simply pass the `CREATE PROPERTY GRAPH` statement to the `execute` method:



Note:

- You can create property graph views using the RDBMS Java API or through SQLcl.
- You can query property graph views using the graph visualization tool or SQLcl.

```
stmt.execute("CREATE PROPERTY GRAPH student_network ...");
```

This results in the creation of the following metadata tables:

```
SQL> SELECT * FROM STUDENT_NETWORK_ELEM_TABLE$;
```

ET_NAME	ET_TYPE	SCHEMA_NAME	TABLE_NAME
PERSON	VERTEX	SCOTT	PERSON
UNIVERSITY	VERTEX	SCOTT	UNIVERSITY
KNOWS	EDGE	SCOTT	KNOWS
STUDENTOF	EDGE	SCOTT	PERSON

```
SQL> SELECT * FROM STUDENT_NETWORK_LABEL$;
```

LABEL_NAME	ET_NAME	ET_TYPE
STUDENT	PERSON	VERTEX
UNIVERSITY	UNIVERSITY	VERTEX
KNOWS	KNOWS	EDGE
STUDENTOF	STUDENTOF	EDGE

```
SQL> SELECT * FROM STUDENT_NETWORK_PROPERTY$;
```

PROPERTY_NAME	ET_NAME	ET_TYPE	LABEL_NAME	COLUMN_NAME
NAME	PERSON	VERTEX	STUDENT	NAME
NAME	UNIVERSITY	VERTEX	UNIVERSITY	NAME

```
SQL> SELECT * FROM STUDENT_NETWORK_KEY$;
```

COLUMN_NAME	COLUMN_NUMBER	KEY_TY	ET_NAME
ID	1	VERTEX	PERSON

```

ID                1 VERTEX UNIVERSITY
PERSON1           1 EDGE   KNOWS
PERSON2           2 EDGE   KNOWS
ID                1 EDGE   STUDENTOF
UNIVERSITY        2 EDGE   STUDENTOF

SQL> SELECT * FROM STUDENT_NETWORK_SRC_DST_KEY$;

ET_NAME      VT_NAME      KEY_TYPE      ET_COLUMN_NAME
ET_COLUMN_NUMBER
-----
KNOWS        PERSON        EDGE_SOURCE   PERSON1
1
KNOWS        PERSON        EDGE_DESTINATION
PERSON2      1
STUDENTOF    PERSON        EDGE_SOURCE   ID
1
STUDENTOF    UNIVERSITY    EDGE_DESTINATION
UNIVERSITY  1

```

You can now run PGQL queries on the property graph view `student_network`.

See [Executing PGQL Queries Against Property Graph Views](#) for more details to create, query and drop property graph views.

- [Loading a Graph into the Graph Server \(PGX\) from a Property Graph View](#)
You can load a graph into the graph server (PGX) from a property graph view by name.
- [Using JSON to Store Vertex and Edge Properties](#)
You can adopt a flexible schema approach in a property graph view by encoding the vertex and edge properties as a single JSON value. You can then map this to a property value in a property graph view.

3.1 Loading a Graph into the Graph Server (PGX) from a Property Graph View

You can load a graph into the graph server (PGX) from a property graph view by name.

You can use the following `PgxSession` method to load a graph from a property graph view by name:

```
readGraphByName(String name, GraphSource source)
```

The arguments used in the method are as follows:

- `name`: Name of the property graph view.
- `source`: Source for the graph. In this case, `PG_VIEW`.

The `readGraphByName(String name, GraphSource source)` method reads the property graph view metadata tables and internally generates the graph configuration to load the graph. You must have `PGX_SESSION_NEW_GRAPH` permission to use this API.

For example you can load the graph from a property graph view as shown:

-
- [JShell](#)
 - [Java](#)
 - [Python](#)

JShell

```
var pgview = session.readGraphByName("bankdataview", GraphSource.PG_VIEW)
$12 ==> PgxGraph[name=bankdataview,N=1000,E=5001,created=1625730942294]
```

Java

```
PgxGraph graph = session.readGraphByName("bankdataview",
GraphSource.PG_VIEW);
Graph: PgxGraph[name=bankdataview,N=1000,E=5001,created=1625732149262]
```

Python

```
>>> graph = session.read_graph_by_name('bankdataview', 'pg_view')
>>> graph
PgxGraph(name: bankdataview, v: 1000, e: 5001, directed: True, memory(Mb): 0)
```


3.2 Using JSON to Store Vertex and Edge Properties

You can adopt a flexible schema approach in a property graph view by encoding the vertex and edge properties as a single JSON value. You can then map this to a property value in a property graph view.

The property graph schema graphs by nature provide schema flexibility. You can add arbitrary labels and properties to the graph without having to issue explicit schema updates (for example, through `ALTER TABLE` statement). Property graph views, on the other hand, do not provide such flexibility by nature since adding a new label requires adding a new vertex or edge table, and adding a new property requires adding a new column, both of which are schema update operations. However, through the use of JSON you can model schema flexibility on top of property graph views.

For example, consider the following graph which represent financial transactions between two `Account` vertices. The `Account` can be owned either by a `Person` or a `Company`.

Figure 3-1 Financial Transactions Graph



You can create a single table for storing all the vertices and another single table for storing all the edges, as shown:

```
CREATE TABLE fin_vertex_table (
  id NUMBER PRIMARY KEY,
  properties VARCHAR2(2000)
);

INSERT INTO fin_vertex_table VALUES ( 1,
'{"type":"Person","name":"Nikita"}');
INSERT INTO fin_vertex_table VALUES ( 2,
'{"type":"Person","name":"Camille"}');
INSERT INTO fin_vertex_table VALUES ( 3,
'{"type":"Person","name":"Liam"}');
INSERT INTO fin_vertex_table VALUES ( 4,
'{"type":"Company","name":"Oracle"}');
INSERT INTO fin_vertex_table VALUES ( 5,
'{"type":"Account","number":10039}');
INSERT INTO fin_vertex_table VALUES ( 6,
'{"type":"Account","number":2090}');
INSERT INTO fin_vertex_table VALUES ( 7,
'{"type":"Account","number":8021}');
INSERT INTO fin_vertex_table VALUES ( 8,
'{"type":"Account","number":1001}');

CREATE TABLE fin_edge_table (
  id NUMBER PRIMARY KEY,
  src NUMBER REFERENCES fin_vertex_table ( id ),
  dst NUMBER REFERENCES fin_vertex_table ( id ),
  properties VARCHAR2(2000)
);

INSERT INTO fin_edge_table VALUES ( 1, 7, 1, '{"type":"owner"}');
INSERT INTO fin_edge_table VALUES ( 2, 5, 2, '{"type":"owner"}');
INSERT INTO fin_edge_table VALUES ( 3, 6, 3, '{"type":"owner"}');
INSERT INTO fin_edge_table VALUES ( 4, 8, 4, '{"type":"owner"}');
```

```

INSERT INTO fin_edge_table VALUES ( 5, 2, 4, '{"type":"worksFor"}');
INSERT INTO fin_edge_table VALUES ( 6, 5, 7,
'{"type":"transaction","amount":1000.00}');
INSERT INTO fin_edge_table VALUES ( 7, 7, 8,
'{"type":"transaction","amount":1500.30}');
INSERT INTO fin_edge_table VALUES ( 8, 7, 8,
'{"type":"transaction","amount":3000.70}');
INSERT INTO fin_edge_table VALUES ( 9, 8, 6,
'{"type":"transaction","amount":9999.50}');
INSERT INTO fin_edge_table VALUES ( 10, 6, 5,
'{"type":"transaction","amount":9900.00}');

```

As seen in the preceding code, each vertex and edge is represented by a single row in the respective tables. The first column is the unique key of the vertex or the edge. The second and third columns of the edge table are its source key and destination key respectively. The last column of the vertex and edge tables encodes all the properties as well as the labels in a JSON object. A JSON is an unordered set of name and value pairs. Here, you can use such pairs to encode the property names and their values as well as the label's value. In case of the label, you can choose an arbitrary name such as "type" or "label". In this example we use "type".

Because all the labels and properties of a vertex or an edge are encoded as a single JSON value, you do not need to update the schema when new labels or properties are added to the graph. Instead, you can add new labels and properties by inserting additional vertices and edges or by updating the JSON value in the underlying table through SQL.

The following two examples demonstrate how you can extract labels and property values from JSON objects for PGQL on RDBMS and PGQL on PGX respectively.

Example 3-2 Extracting JSON properties using JSON_VALUE (PGQL on RDBMS)

The following code creates a property graph view using the `fin_vertex_table` and `fin_edge_table` tables and executes a PGQL `SELECT` query:

```

PgqlStatement pgqlStmnt = pgqlConn.createStatement();

/* Create the property graph */
pgqlStmnt.execute(
    "CREATE PROPERTY GRAPH financial_transactions " +
    " VERTEX TABLES ( " +
    "   fin_vertex_table PROPERTIES ( properties ) ) " +
    " EDGE TABLES ( " +
    "   fin_edge_table " +
    "     SOURCE KEY ( src ) REFERENCES fin_vertex_table " +
    "     DESTINATION KEY ( dst ) REFERENCES fin_vertex_table " +
    "     PROPERTIES ( properties ) ) " +
    " OPTIONS ( PG_VIEW )");

/* Set the name of the graph so that we can omit the ON clause from queries
*/
pgqlConn.setGraph("FINANCIAL_TRANSACTIONS");

/* PGQL query: find all outgoing transactions from account 8021. Output the
transaction amount and the destination account number. */
PgqlResultSet rs = pgqlStmnt.executeQuery(
    "SELECT JSON_VALUE(trans.properties, '$.amount') AS transaction_amount, " +

```

```

"      JSON_VALUE(account2.properties, '$.number') AS
account_number " +
"FROM MATCH (account1) -[trans]-> (account2) " +
"WHERE JSON_VALUE(account1.properties, '$.number') = 8021 " +
" AND JSON_VALUE(trans.properties, '$.type') = 'transaction'");

rs.print();
rs.close();
pgqlStmt.close();

```

In the preceding code, the `CREATE PROPERTY GRAPH` statement maps the JSON column into a property named "properties". This property will thus contain all the labels and properties of the vertex or the edge. The PGQL `SELECT` query extracts these labels and properties using `JSON_VALUE`.

For example, instead of `account1.number = 8021`, you must use `JSON_VALUE(account1.properties, '$.number') = 8021`. This causes the query to become a bit lengthier.

The output of the Java code is:

```

+-----+
| AMOUNT | ACCOUNT_NUMBER |
+-----+
| 1500.3 | 1001           |
| 3000.7 | 1001           |
+-----+

```

Example 3-3 Using a UDF to extract a JSON property value (PGQL on PGX)

This example consists of two parts. The first part shows the creation of a UDF and the second part shows loading of the graph into the graph server (PGX) followed by the execution of a PGQL query using the UDF.

Since the Graph Server (PGX) does not have a built-in `JSON_VALUE` function like in PGQL on RDBMS, you can create a Java UDF instead.

Create the Java class (`MyJsonUtils.java`) that implements the UDF:

```

import com.fasterxml.jackson.core.JsonProcessingException;
import com.fasterxml.jackson.databind.JsonNode;
import com.fasterxml.jackson.databind.ObjectMapper;

public class MyJsonUtils {

    private final static ObjectMapper mapper = new ObjectMapper();

    public static String get_prop(String json_string, String prop_name)
    throws JsonProcessingException {
        JsonNode node = mapper.readTree(json_string);
        return node.path(prop_name).asText();
    }
}

```


Compile the class with the JARs from `/opt/oracle/graph/pgx/server/lib/*` added to the class path. This is because the library folder contains the necessary Jackson libraries that are required to parse the JSON.

```
mkdir ./target
javac -classpath ./opt/oracle/graph/pgx/server/lib/* -d ./target *.java
cd target
jar cvf MyJsonUtils.jar *
```

Using the following UDF JSON configuration file (`my_udfs.json`), you can now register the Java UDF on the graph server (PGX) by following step-3 to step-6 in [User-Defined Functions \(UDFs\) in PGX](#):

```
{
  "user_defined_functions": [
    {
      "namespace": "my",
      "function_name": "get_prop",
      "language": "java",
      "implementation_reference": "MyJsonUtils",
      "return_type": "string",
      "arguments": [
        {
          "name": "json_string",
          "type": "string"
        },
        {
          "name": "prop_name",
          "type": "string"
        }
      ]
    }
  ]
}
```

On implementing the UDF for extracting property values from the JSON, you can now load the graph into the Graph Server (PGX) and issue a PGQL query:

```
/* Load the graph into the Graph Server (PGX) */
ServerInstance instance = GraphServer.getInstance("http://localhost:7007",
username, password.toCharArray());
session = instance.createSession("my-session");
PgxGraph g = session.readGraphByName("FINANCIAL_TRANSACTIONS",
GraphSource.PG_VIEW);

/* PGQL query: find all shortest paths from account 10039 to account 2090
following only outgoing transaction
edges. Output the list of transaction amounts along each path as well as
the total amount of the transactions
along each path. */
g.queryPgql(
" SELECT LISTAGG(my.get_prop(e.properties, 'amount'), ' + ') || ' = ' AS
amounts_along_path, " +
```

```

"          SUM(CAST(my.get_prop(e.properties, 'amount') AS DOUBLE))
AS total_amount " +
"    FROM MATCH ALL SHORTEST (a) (-[e]-> WHERE
my.get_prop(e.properties, 'type') = 'transaction')* (b) " +
"    WHERE my.get_prop(a.properties, 'number') = '10039' AND " +
"          my.get_prop(b.properties, 'number') = '2090' " +
"ORDER BY total_amount").print().close();

```

The output of the Java code is:

```

+-----+
| amounts_along_path          | total_amount |
+-----+
| 1000.0 + 1500.3 + 9999.5 =  | 12499.8     |
| 1000.0 + 3000.7 + 9999.5 =  | 14000.2     |
+-----+

```

4

Using the In-Memory Graph Server (PGX)

The in-memory Graph server of Oracle Graph supports a set of analytical functions.

This chapter provides examples using the in-memory Graph Server (also referred to as Property Graph In-Memory Analytics, and often abbreviated as PGX in the Javadoc, command line, path descriptions, error messages, and examples). It contains the following major topics.

- [Overview of the In-Memory Graph Server \(PGX\)](#)
The In-Memory Graph Server (PGX) is an in-memory accelerator for fast, parallel graph query and analytics. The server uses light-weight in-memory data structures to enable fast execution of graph algorithms.
- [User Authentication and Authorization](#)
The Oracle Graph server (PGX) uses an Oracle Database as identity manager. Both username and password based as well as Kerberos based authentication is supported.
- [About Vertex and Edge IDs](#)
- [Reading Graphs from Oracle Database into the Graph Server \(PGX\)](#)
Once logged in to the graph server (PGX), you can now read graphs from the database into the graph server without specifying any connection information in the graph configuration.
- [Keeping the Graph in Oracle Database Synchronized with the Graph Server](#)
You can use the `FlashbackSynchronizer` API to automatically apply changes made to graph in the database to the corresponding `PgxGraph` object in memory, thus keeping both synchronized.
- [Optimizing Graphs for Read Versus Updates in the In-Memory Graph Server \(PGX\)](#)
- [Storing a Graph Snapshot on Disk](#)
After reading a graph into memory using either Java or the Shell, if you make some changes to the graph such as running the PageRank algorithm and storing the values as vertex properties, you can store this snapshot of the graph on disk.
- [Executing Built-in Algorithms](#)
The in-memory graph server (PGX) contains a set of built-in algorithms that are available as Java APIs.
- [Using Custom PGX Graph Algorithms](#)
A custom PGX graph algorithm allows you to write a graph algorithm in Java syntax and have it automatically compiled to an efficient parallel implementation.
- [Creating Subgraphs](#)
You can create subgraphs based on a graph that has been loaded into memory. You can use filter expressions or create bipartite subgraphs based on a vertex (node) collection that specifies the left set of the bipartite graph.
- [Using Automatic Delta Refresh to Handle Database Changes](#)
You can automatically refresh (auto-refresh) graphs periodically to keep the in-memory graph synchronized with changes to the property graph stored in the property graph tables in Oracle Database (VT\$ and GE\$ tables).

- [Starting the In-Memory Graph Server \(PGX\)](#)
This section describes the commands to start and stop the in-memory graph server(PGX).
- [Connecting to the In-Memory Graph Server \(PGX\)](#)
This section explains how to connect to the in-memory graph server (PGX) running in remote mode or when deployed as a web application on Apache Tomcat or Oracle WebLogic Server.
- [Using Graph Server \(PGX\) as a Library](#)
- [User-Defined Functions \(UDFs\) in PGX](#)
User-defined functions (UDFs) allow users of PGX to add custom logic to their PGQL queries or custom graph algorithms, to complement built-in functions with custom requirements.
- [Using HAProxy for PGX Load Balancing and High Availability](#)
HAProxy is a high-performance TCP/HTTP load balancer and proxy server that allows multiplexing incoming requests across multiple web servers.

4.1 Overview of the In-Memory Graph Server (PGX)

The In-Memory Graph Server (PGX) is an in-memory accelerator for fast, parallel graph query and analytics. The server uses light-weight in-memory data structures to enable fast execution of graph algorithms.

There are multiple options to load a graph into the graph server either from Oracle Database or from files.

The graph server can be deployed standalone (it includes an embedded Apache Tomcat instance), or deployed in Oracle WebLogic Server or Apache Tomcat.

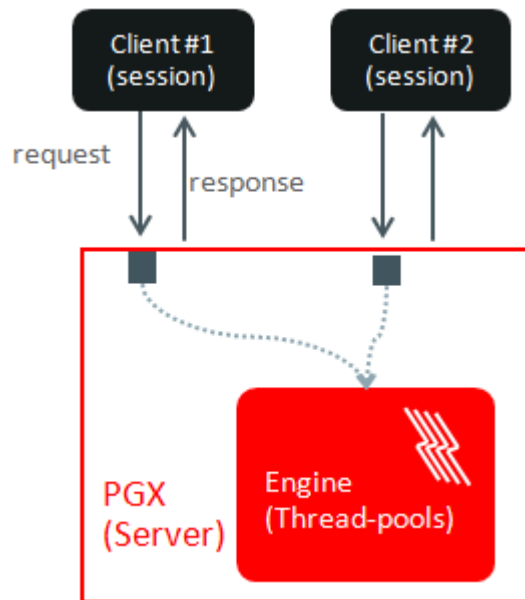
- [Design of the In-Memory Graph Server \(PGX\)](#)
- [Usage Modes of the In-memory Graph Server \(PGX\)](#)

4.1.1 Design of the In-Memory Graph Server (PGX)

The design of the in-memory graph server (PGX) is based on a Server-Client usage model. See [Usage Modes of the In-memory Graph Server \(PGX\)](#) for more details on the different graph server (PGX) execution modes.

The following figure shows the graph server (PGX) design:

Figure 4-1 Graph Server (PGX) Design



The core concepts of the graph server (PGX) design are as follows:

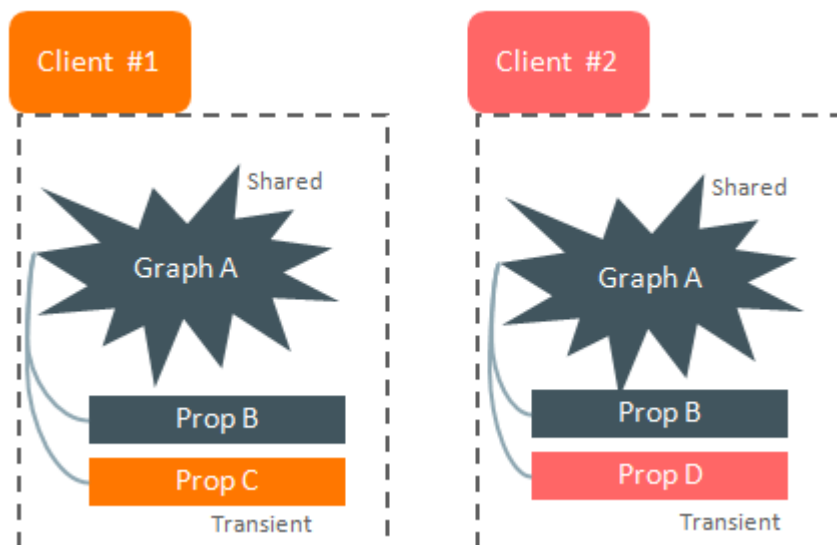
- Multiple graph clients can connect to the in-memory graph server at the same time.
- Each client request is processed by the graph server asynchronously. The client requests are queued up first and processed later, when resources are available. The client can poll the server to check if a request has been finished.
- Internally, the server maintains its own engine (thread pools) for running parallel graph algorithms and queries. The engine tries to process each analytics request concurrently with as many threads as possible.

Isolation Between Concurrent Clients

The graph server (PGX) supports data isolation between concurrent clients. Each client has its own private workspace, called session. Sessions are isolated from each other. Each client can load a graph instance into its own session, independently from other clients.

If multiple clients load the same graph instance the graph server can share one graph instance between multiple clients under the hood. Each client can add additional vertex or edge properties to a loaded graph in its own session. Such properties are transient properties, private to each session and not visible to another session as shown in the following figure:

Figure 4-2 Session and Transient Properties



Similarly, if a client creates a mutated version of the loaded graph, the graph server will create a private graph instance for that client.

4.1.2 Usage Modes of the In-memory Graph Server (PGX)

This section presents an overview of the different usage modes of the graph server (PGX). The graph server can be executed in one of the following usage modes.

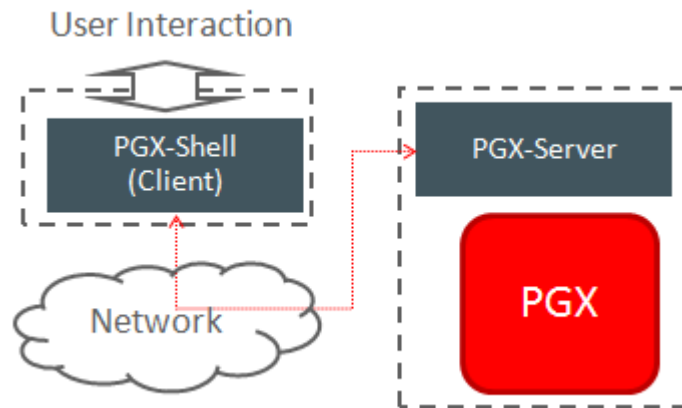
Remote Server Mode

In the remote server mode, the main PGX execution engine is deployed as a RESTful application on a powerful server machine, and you can connect to it remotely from your machine using graph shell. Also, multiple clients can connect to the same graph server (PGX) at the same time and therefore the graph server is time-shared among these clients.

See [Interactive Graph Shell CLIs](#) for more information on the graph shell.

The following figure shows the graph server (PGX) in a remote execution mode:

Figure 4-3 Remote Server Mode



The remote server mode is useful for the following situations where you want to:

- Perform graph analysis on a large data set with a powerful server-class machine that has many cores and a large memory.
- The server-class machine is shared by multiple clients.

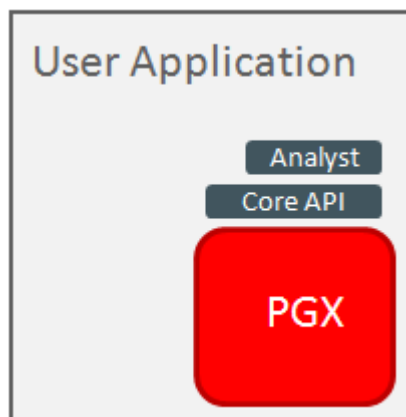
See [Starting the In-Memory Graph Server \(PGX\)](#) for instructions on how to start the graph server (PGX) in remote server mode.

Using Graph Server (PGX) as a Library

You can also include the graph server (PGX) as a normal Java library in your application.

The following figure shows the graph server (PGX) used as a library in an application:

Figure 4-4 PGX as a Library



The embedded mode is useful when you want to build an application having graph analysis as a part of its functionality.

See [Using Graph Server \(PGX\) as a Library](#) for more information.

Deploying Graph Server (PGX) as Servlet Web Application

You can deploy the graph server (PGX) as a web application using Apache Tomcat or Oracle WebLogic Server.

See [Deploying Oracle Graph Server to a Web Server](#) for instructions to deploy the graph server (PGX) in Apache Tomcat or Oracle WebLogic Server.

4.2 User Authentication and Authorization

The Oracle Graph server (PGX) uses an Oracle Database as identity manager. Both username and password based as well as Kerberos based authentication is supported.

The actions that you are allowed to do on the graph server are determined by the privileges enabled by roles that have been granted to you in the Oracle Database.

- [Privileges and Roles in Oracle Database](#)
All database users that work with graphs require the `CREATE SESSION` privilege in the database.
- [Basic Steps for Using an Oracle Database for Authentication](#)
You can follow the steps explained in this section to authenticate users to the graph server (PGX).
- [Prepare the Graph Server for Database Authentication](#)
Locate the `pgx.conf` file of your installation.
- [Store the Database Password in a Keystore](#)
- [Token Expiration](#)
By default, tokens are valid for 1 hour.
- [Advanced Access Configuration](#)
You can customize the following fields inside the `pgx_realm` block in the `pgx.conf` file to customize login behavior.
- [Customizing Roles and Permissions](#)
You can fully customize the permissions to roles mapping by adding and removing roles and specifying permissions for a role. You can also authorize individual users instead of roles.
- [Revoking Access to the Graph Server](#)
To revoke a user's ability to access the graph server, either drop the user from the database or revoke the corresponding roles from the user, depending on how you defined the access rules in your `pgx.conf` file.
- [Examples of Custom Authorization Rules](#)
You can define custom authorization rules for developers.
- [Kerberos Enabled Authentication](#)
The graph server (PGX) can authenticate users using an Oracle Database with Kerberos enabled as identity provider.

4.2.1 Privileges and Roles in Oracle Database

All database users that work with graphs require the `CREATE SESSION` privilege in the database.

Roles that are created for working with graphs are in [Table 4-1](#). These roles are created when you install the PL/SQL package of the Oracle Graph Server and Client distribution on the target database.

Table 4-1 Privileges and Roles in Oracle Database

Role	Operations enabled by this role	Used By
PGX_SESSION_CREATE	Create a new PGX session using the <code>ServerInstance.createSession</code> API.	Graph developers and graph users
PGX_SERVER_GET_INFO	Get status information on the PGX instance using the Admin API .	Users who administer PGX
PGX_SERVER_MANAGE (includes PGX_SERVER_GET_INFO)	Manage the PGX instance using the Admin API to stop or restart PGX.	Users who administer PGX
PGX_SESSION_NEW_GRAPH	Create a new graph in PGX by loading from the database using a config file, using the <code>CREATE PROPERTY GRAPH</code> statement in PGQL, creating a sub-graph from another graph, or using the <code>GraphBuilder</code> .	Graph developers and graph users
PGX_SESSION_GET_PUBLISHED_GRAPH	Query and view graphs published by another user to the public namespace.	Graph developers and graph users
PGX_SESSION_ADD_PUBLISHED_GRAPH (includes PGX_SESSION_GET_PUBLISHED_GRAPH)	Publish a graph to the public namespace.	Graph developers
PGX_SESSION_COMPILE_ALGORITHM	Compile an algorithm using the PGX <code>Algorithm API</code> .	Graph developers
PGX_SESSION_READ_MODEL	Load and use an ML model using <code>PgxML</code> .	Graph developers
PGX_SESSION_MODIFY_MODEL	Create, train, and store an ML model using <code>PgxML</code> .	Graph developers

Few additional roles are also created to group multiple roles together. They provide a convenient way to grant multiple roles to database users. See [Mapping Graph Server Roles to Default Privileges](#) for more information on these additional roles.

You can create additional groups that are useful for your application, as described in [Adding and Removing Roles](#) and [Defining Permissions for Individual Users](#).

4.2.2 Basic Steps for Using an Oracle Database for Authentication

You can follow the steps explained in this section to authenticate users to the graph server (PGX).

1. Use an Oracle Database version that is supported by Oracle Graph Server and Client: version 12.2 or later, including Autonomous Database.

2. Be sure that you have ADMIN access (or SYSDBA access for non-autonomous databases) to grant and revoke users access to the graph server (PGX).
3. Be sure that all existing users to which you plan to grant access to the graph server have at least the CREATE SESSION privilege granted.
4. Be sure that the database is accessible via JDBC from the host where the Graph Server runs.
5. As ADMIN (or SYSDBA on non-autonomous databases), run the following procedure to create the roles required by the graph server:

 **Note:**

You can skip this step if you install the PL/SQL packages as part of the Oracle Graph Server and Client installation. All the roles shown in the following code are created as part of the PL/SQL installation automatically. You need to add them separately only if you are using Oracle Graph Server and Client with Autonomous Database. You can run this code using Database Actions in Oracle Cloud Infrastructure Console.

```

DECLARE
  PRAGMA AUTONOMOUS_TRANSACTION;
  role_exists EXCEPTION;
  PRAGMA EXCEPTION_INIT(role_exists, -01921);
  TYPE graph_roles_table IS TABLE OF VARCHAR2(50);
  graph_roles graph_roles_table;
BEGIN
  graph_roles := graph_roles_table(
    'GRAPH_DEVELOPER',
    'GRAPH_ADMINISTRATOR',
    'GRAPH_USER',
    'PGX_SESSION_CREATE',
    'PGX_SERVER_GET_INFO',
    'PGX_SERVER_MANAGE',
    'PGX_SESSION_READ_MODEL',
    'PGX_SESSION_MODIFY_MODEL',
    'PGX_SESSION_NEW_GRAPH',
    'PGX_SESSION_GET_PUBLISHED_GRAPH',
    'PGX_SESSION_COMPILE_ALGORITHM',
    'PGX_SESSION_ADD_PUBLISHED_GRAPH');
  FOR elem IN 1 .. graph_roles.count LOOP
    BEGIN
      dbms_output.put_line('create_graph_roles: ' || elem || ');
    CREATE ROLE ' || graph_roles(elem));
      EXECUTE IMMEDIATE 'CREATE ROLE ' || graph_roles(elem);
    EXCEPTION
      WHEN role_exists THEN
        dbms_output.put_line('create_graph_roles: role already
exists. continue');
      WHEN OTHERS THEN
        RAISE;
    
```

```

        END;
    END LOOP;
EXCEPTION
    when others then
        dbms_output.put_line('create_graph_roles: hit error ');
        raise;
END;
/

```

6. Assign default permissions to the roles GRAPH_DEVELOPER, GRAPH_USER and GRAPH_ADMINISTRATOR to group multiple permissions together.

 **Note:**

You can skip this step if you install the PL/SQL packages as part of the Oracle Graph Server and Client installation. All the grants shown in the following code are executed as part of the PL/SQL installation automatically. You need to execute these grants separately only if you are using Oracle Graph Server and Client with Autonomous Database. You can run this code using Database Actions in Oracle Cloud Infrastructure Console.

```

GRANT PGX_SESSION_CREATE TO GRAPH_ADMINISTRATOR;
GRANT PGX_SERVER_GET_INFO TO GRAPH_ADMINISTRATOR;
GRANT PGX_SERVER_MANAGE TO GRAPH_ADMINISTRATOR;
GRANT PGX_SESSION_CREATE TO GRAPH_DEVELOPER;
GRANT PGX_SESSION_NEW_GRAPH TO GRAPH_DEVELOPER;
GRANT PGX_SESSION_GET_PUBLISHED_GRAPH TO GRAPH_DEVELOPER;
GRANT PGX_SESSION_MODIFY_MODEL TO GRAPH_DEVELOPER;
GRANT PGX_SESSION_READ_MODEL TO GRAPH_DEVELOPER;
GRANT PGX_SESSION_CREATE TO GRAPH_USER;
GRANT PGX_SESSION_GET_PUBLISHED_GRAPH TO GRAPH_USER;

```

7. Assign roles to all the database developers who should have access to the graph server (PGX). For example:

```
GRANT graph_developer TO <graphuser>
```

where <graphuser> is a user in the database. You can also assign individual permissions (roles prefixed with PGX_) to users directly.

8. Assign the administrator role to users who should have administrative access. For example:

```
GRANT graph_administrator to <administratoruser>
```

where <administratoruser> is a user in the database.

4.2.3 Prepare the Graph Server for Database Authentication

Locate the `pgx.conf` file of your installation.

If you installed the graph server via RPM, the file is located at: `/etc/oracle/graph/pgx.conf`

If you use the `webapps` package to deploy into Tomcat or WebLogic Server, the `pgx.conf` file is located inside the web application archive file (WAR file) at: `WEB-INF/classes/pgx.conf`

Tip: On Linux, you can use `vim` to edit the file directly inside the WAR file without unzipping it first. For example:

```
vim graph-server-<version>-pgx<version>.war
```

Inside the `pgx.conf` file, locate the `jdbc_url` line of the realm options:

```
...
"pgx_realm": {
  "implementation": "oracle.pg.identity.DatabaseRealm",
  "options": {
    "jdbc_url": "<REPLACE-WITH-DATABASE-URL-TO-USE-FOR-
AUTHENTICATION>",
    "token_expiration_seconds": 3600,
  }
}
...
```

Replace the text with the JDBC URL pointing to your database that you configured in the previous step. For example:

```
...
"pgx_realm": {
  "implementation": "oracle.pg.identity.DatabaseRealm",
  "options": {
    "jdbc_url": "jdbc:oracle:thin:@myhost:1521/mysevice",
    "token_expiration_seconds": 3600,
  }
}
...
```

If you are using an Autonomous Database, specify the JDBC URL like this:

```
...
"pgx_realm": {
  "implementation": "oracle.pg.identity.DatabaseRealm",
  "options": {
    "jdbc_url": "jdbc:oracle:thin:@my_identifier_low?TNS_ADMIN=/etc/
oracle/graph/wallet",
    "token_expiration_seconds": 3600,
  }
}
...
```

where `/etc/oracle/graph/wallet` is an example path to the unzipped wallet file that you downloaded from your Autonomous Database service console, and `my_identifier_low` is one of the connect identifiers specified in `/etc/oracle/graph/wallet/tnsnames.ora`.

Now, start the graph server. If you installed via RPM, execute the following command as a root user or with `sudo`:

```
sudo systemctl start pgx
```

4.2.4 Store the Database Password in a Keystore

PGX requires a database account to read data from the database into memory. The account should be a low-privilege account (see [Security Best Practices with Graph Data](#)).

As described in [Reading Graphs from Oracle Database into the Graph Server \(PGX\)](#), you can read data from the database into the graph server without specifying additional authentication as long as the token is valid for that database user. But if you want to access a graph from a different user, you can do so, as long as that user's password is stored in a Java Keystore file for protection.

You can use the `keytool` command that is bundled together with the JDK to generate such a keystore file on the command line. See the following script as an example:

```
# Add a password for the 'database1' connection
keytool -importpass -alias database1 -keystore keystore.p12
# 1. Enter the password for the keystore
# 2. Enter the password for the database

# Add another password (for the 'database2' connection)
keytool -importpass -alias database2 -keystore keystore.p12

# List what's in the keystore using the keytool
keytool -list -keystore keystore.p12
```

If you are using Java version 8 or lower, you should pass the additional parameter `-storetype pkcs12` to the `keytool` commands in the preceding example.

You can store more than one password into a single keystore file. Each password can be referenced using the alias name provided.

- [Either, Write the PGX graph configuration file to load from the property graph schema](#)
- [Or, Write the PGX graph configuration file to load a graph directly from relational tables](#)
- [Read the data](#)
- [Secure coding tips for graph client applications](#)

Either, Write the PGX graph configuration file to load from the property graph schema

Next write a PGX graph configuration file in JSON format. The file tells PGX where to load the data from, how the data looks like and the keystore alias to use. The following example shows a graph configuration to read data stored in the Oracle property graph format.

```
{
  "format": "pg",
  "db_engine": "rdbms",
  "name": "hr",
  "jdbc_url": "jdbc:oracle:thin:@myhost:1521/orcl",
  "username": "hr",
```

```

"keystore_alias": "database1",
"vertex_props": [{
  "name": "COUNTRY_NAME",
  "type": "string"
}, {
  "name": "DEPARTMENT_NAME",
  "type": "string"
}, {
  "name": "SALARY",
  "type": "double"
}],
"partition_while_loading": "by_label",
"loading": {
  "load_vertex_labels": true,
  "load_edge_label": true
}
}

```

(For the full list of available configuration fields, including their meanings and default values, see [Graph Configuration Options](#).)

Or, Write the PGX graph configuration file to load a graph directly from relational tables

The following example loads a subset of the HR sample data from relational tables directly into PGX as a graph. The configuration file specifies a mapping from relational to graph format by using the concept of vertex and edge providers.



Note:

Specifying the `vertex_providers` and `edge_providers` properties loads the data into an optimized representation of the graph.

```

{
  "name": "hr",
  "jdbc_url": "jdbc:oracle:thin:@myhost:1521/orcl",
  "username": "hr",
  "keystore_alias": "database1",
  "vertex_id_strategy": "no_ids",
  "vertex_providers": [
    {
      "name": "Employees",
      "format": "rdbms",
      "database_table_name": "EMPLOYEES",
      "key_column": "EMPLOYEE_ID",
      "key_type": "string",
      "props": [
        {
          "name": "FIRST_NAME",
          "type": "string"
        }
      ]
    }
  ]
}

```

```
        "name": "LAST_NAME",
        "type": "string"
    },
    {
        "name": "EMAIL",
        "type": "string"
    },
    {
        "name": "SALARY",
        "type": "long"
    }
]
},
{
    "name": "Jobs",
    "format": "rdbms",
    "database_table_name": "JOBS",
    "key_column": "JOB_ID",
    "key_type": "string",
    "props": [
        {
            "name": "JOB_TITLE",
            "type": "string"
        }
    ]
},
{
    "name": "Departments",
    "format": "rdbms",
    "database_table_name": "DEPARTMENTS",
    "key_column": "DEPARTMENT_ID",
    "key_type": "string",
    "props": [
        {
            "name": "DEPARTMENT_NAME",
            "type": "string"
        }
    ]
}
],
"edge_providers": [
    {
        "name": "WorksFor",
        "format": "rdbms",
        "database_table_name": "EMPLOYEES",
        "key_column": "EMPLOYEE_ID",
        "source_column": "EMPLOYEE_ID",
        "destination_column": "EMPLOYEE_ID",
        "source_vertex_provider": "Employees",
        "destination_vertex_provider": "Employees"
    },
    {
        "name": "WorksAs",
        "format": "rdbms",
        "database_table_name": "EMPLOYEES",
```

```

        "key_column": "EMPLOYEE_ID",
        "source_column": "EMPLOYEE_ID",
        "destination_column": "JOB_ID",
        "source_vertex_provider": "Employees",
        "destination_vertex_provider": "Jobs"
    },
    {
        "name": "WorkedAt",
        "format": "rdbms",
        "database_table_name": "JOB_HISTORY",
        "key_column": "EMPLOYEE_ID",
        "source_column": "EMPLOYEE_ID",
        "destination_column": "DEPARTMENT_ID",
        "source_vertex_provider": "Employees",
        "destination_vertex_provider": "Departments",
        "props": [
            {
                "name": "START_DATE",
                "type": "local_date"
            },
            {
                "name": "END_DATE",
                "type": "local_date"
            }
        ]
    }
]
}

```

Read the data

Now you can instruct PGX to connect to the database and read the data by passing in both the keystore and the configuration file to PGX, using one of the following approaches:

- **Interactively in the graph shell**
If you are using the graph shell, start it with the `--secret_store` option. It will prompt you for the keystore password and then attach the keystore to your current session. For example:

```

cd /opt/oracle/graph
./bin/opg4j --secret_store /etc/my-secrets/keystore.p12

enter password for keystore /etc/my-secrets/keystore.p12:

```

Inside the shell, you can then use normal PGX APIs to read the graph into memory by passing the JSON file you just wrote into the `readGraphWithProperties` API:

```

opg4j> var graph = session.readGraphWithProperties("config.json")
graph ==> PgxGraph[name=hr,N=215,E=415,created=1576882388130]

```

- **As a PGX preloaded graph**

As a server administrator, you can instruct PGX to load graphs into memory upon server startup. To do so, modify the PGX configuration file at `/etc/oracle/graph/pgx.conf` and add the path the graph configuration file to the `preload_graphs` section. For example:

```
{
  ...
  "preload_graphs": [{
    "name": "hr",
    "path": "/path/to/config.json"
  }],
  "authorization": [{
    "pgx_role": "GRAPH_DEVELOPER",
    "pgx_permissions": [{
      "preloaded_graph": "hr",
      "grant": "read"
    }]
  }],
  ....
}]
}
```

As root user, edit the service file at `/etc/systemd/system/pgx.service` and change the `ExecStart` command to specify the location of the keystore containing the password:

```
ExecStart=/bin/bash start-server --secret-store /etc/keystore.p12
```

 **Note:**

Please note that `/etc/keystore.p12` must not be password protected for this to work. Instead protect the file via file system permission that is only readable by `oraclegraph` user.

After the file is edited, reload the changes using:

```
sudo systemctl daemon-reload
```

Finally start the server:

```
sudo systemctl start pgx
```

- **In a Java application**

To register a keystore in a Java application, use the `registerKeystore()` API on the `PgxSession` object. For example:

```
import oracle.pgx.api.*;

class Main {

    public static void main(String[] args) throws Exception {
```

```
String baseUrl = args[0];
String keystorePath = "/etc/my-secrets/keystore.p12";
char[] keystorePassword = args[1].toCharArray();
String graphConfigPath = args[2];
ServerInstance instance = Pgx.getInstance(baseUrl);
try (PgxSession session = instance.createSession("my-session"))
{
    session.registerKeystore(keystorePath, keystorePassword);
    PgxGraph graph =
session.readGraphWithProperties(graphConfigPath);
    System.out.println("N = " + graph.getNumVertices() + " E = "
+ graph.getNumEdges());
}
}
```

You can compile and run the preceding sample program using the Oracle Graph Client package. For example:

```
cd $GRAPH_CLIENT
// create Main.java with above contents
javac -cp 'lib/*' Main.java
java -cp '.:conf:lib/*' Main http://myhost:7007 MyKeystorePassword
path/to/config.json
```

Secure coding tips for graph client applications

When writing graph client applications, make sure to never store any passwords or other secrets in clear text in any files or in any of your code.

Do not accept passwords or other secrets through command line arguments either. Instead, use `Console.html#readPassword()` from the JDK.

4.2.5 Token Expiration

By default, tokens are valid for 1 hour.

Internally, the graph client automatically renews tokens which are about to expire in less than 30 minutes. This is also configurable by re-authenticating your credentials with the database. By default, tokens can only be automatically renewed for up to 24 times, then you need to login again.

If the maximum amount of auto-renewals is reached, you can log in again without losing any of your session data by using the `GraphServer#reauthenticate(instance, "<user>", "<password>")` API.

Note:

If a session time out occurs before you re-authenticate, then you may lose your session data.

For example:

```
opg4j> var graph = session.readGraphWithProperties(config) // fails because
token cannot be renewed anymore
opg4j> GraphServer.reauthenticate(instance, "<user>", "<password>") // log
in again
opg4j> var graph = session.readGraphWithProperties(config) // works
now
```

4.2.6 Advanced Access Configuration

You can customize the following fields inside the `pgx_realm` block in the `pgx.conf` file to customize login behavior.

Table 4-2 Advanced Access Configuration Options

Field Name	Explanation	Default
<code>token_expiration_seconds</code>	After how many seconds the generated bearer token will expire.	3600 (1 hour)
<code>connect_timeout_milliseconds</code>	After how many milliseconds an connection attempt to the specified JDBC URL will time out, resulting in the login attempt being rejected.	10000
<code>max_pool_size</code>	Maximum number of JDBC connections allowed per user. If the number is reached, attempts to read from the database will fail for the current user.	64
<code>max_num_users</code>	Maximum number of active, signed in users to allow. If this number is reached, the graph server will reject login attempts.	512
<code>max_num_token_refresh</code>	Maximum amount of times a token can be automatically refreshed before requiring a login again.	24

To configure the refresh time on the client side before token expiration, use the following API to login:

```
int refreshTimeBeforeTokenExpiry = 900; // in seconds, default is 1800 (30 minutes)
ServerInstance instance = GraphServer.getInstance("https://localhost:7007", "<database
user>", "<database password>",
refreshTimeBeforeTokenExpiry);
```

**Note:**

The preceding options work only if the realm implementation is configured to be `oracle.pg.identity.DatabaseRealm`.

4.2.7 Customizing Roles and Permissions

You can fully customize the permissions to roles mapping by adding and removing roles and specifying permissions for a role. You can also authorize individual users instead of roles.

This topic includes examples of how to customize the permission mapping.

- [Checking Graph Permissions Using API](#)
- [Adding and Removing Roles](#)
You can add new role permission mappings or remove existing mappings by modifying the authorization list.
- [Defining Permissions for Individual Users](#)
In addition to defining permissions for roles, you can define permissions for individual users.
- [Defining Permissions to Use Custom Graph Algorithms](#)
You can define permissions to allow developers to compile custom graph algorithms.

4.2.7.1 Checking Graph Permissions Using API

You can view your roles and graph permissions using the following PGX API methods:

Table 4-3 API for Checking Graph Permissions

Class	Method	Description
ServerInstance	<code>getPgxUsername()</code>	Name of the current user
ServerInstance	<code>getPgxUserRoles()</code>	Role names of the current user
ServerInstance	<code>getPgxGenericPermissions()</code>	Non-graph (system) permissions of the current user: <ul style="list-style-type: none"> • Pgx system permissions • File-location permissions
PgxGraph	<code>getPermission()</code>	Permission on the graph instance for a current user

You can get all permission-related information using the API in JShell as shown:

- [JShell](#)
- [Java](#)

JShell

```

/bin/opg4j -b "https://<host>:<port>" -u "<graphuser>"
opg4j> instance
instance ==> ServerInstance[embedded=false,baseUrl=https://
<host>:<port>,serverVersion=null]
opg4j>instance.getPgxUsername()
$2 ==> "ORACLE"
opg4j>instance.getPgxUserRoles()
$3 ==> [GRAPH_DEVELOPER]
opg4j>instance.getPgxGenericPermissions()
$4 ==> [PGX_SESSION_CREATE, PGX_SESSION_READ_MODEL,
PGX_SESSION_ADD_PUBLISHED_GRAPH, PGX_SESSION_NEW_GRAPH,
PGX_SESSION_GET_PUBLISHED_GRAPH, PGX_SESSION_MODIFY_MODEL]
opg4j>var g = session.readGraphWithProperties("bank_graph_analytics.json")
g ==> PgxGraph[name=bank_graph_analytics,N=1000,E=5001,created=1625697341555]
opg4j>g.getPermission() // To get graph permissions
$9 ==> MANAGE

```

Java

```

import oracle.pg.rdbms.*;
import java.sql.Connection;
import java.sql.Statement;
import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlStatement;
import oracle.pgx.api.*;
import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;
import java.nio.file.Files;
import java.nio.file.Path;

/**
 * This example shows how to get all permissions.
 */
public class GetPermissions
{

    public static void main(String[] args) throws Exception
    {
        int idx=0;
        String host          = args[idx++];
        String port          = args[idx++];
        String sid           = args[idx++];
        String user          = args[idx++];
        String password      = args[idx++];
        String graph         = args[idx++];

        Connection conn = null;
        PgxPreparedStatement stmt = null;

        try {

            // Get a jdbc connection

```

```
PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
pds.setURL("jdbc:oracle:thin:@"+host+": "+port + "/" + sid);
pds.setUser(user);
pds.setPassword(password);
conn = pds.getConnection();
conn.setAutoCommit(false);

ServerInstance instance = GraphServer.getInstance("http://
localhost:7007", user, password.toCharArray());
PgxSession session = instance.createSession("my-session");

var statement = Files.readString(Path.of("/media/sf_Linux/Java/
create-pg.pgql"));
stmt = session.preparePgql(statement);
stmt.execute();

PgxGraph g = session.getGraph(graph);
System.out.println("Graph: " + g);

String userName = instance.getPgxUsername();
var userRoles = instance.getPgxUserRoles();
var genericPermissions = instance.getPgxGenericPermissions();
String graphPermission = g.getPermission().toString();

System.out.println("Username is " + userName);
System.out.println("User Roles are " + userRoles);
System.out.println("Generic permissions are " +
genericPermissions);
System.out.println("Graph permission is " + graphPermission);

}

finally {
    // close the sql statment
    if (stmt != null) {
        stmt.close();
    }
    // close the connection
    if (conn != null) {
        conn.close();
    }
}
}
```

On execution, the code gives the following output:

```
Graph: PgxGraph[name=BANK_GRAPH_PG,N=1000,E=5001,created=1625731370402]
Username is ORACLE
User Roles are [GRAPH_DEVELOPER]
Generic permissions are [PGX_SESSION_MODIFY_MODEL, PGX_SESSION_CREATE,
PGX_SESSION_NEW_GRAPH, PGX_SESSION_READ_MODEL,
```

```
PGX_SESSION_ADD_PUBLISHED_GRAPH, PGX_SESSION_GET_PUBLISHED_GRAPH]
Graph permission is MANAGE
```

4.2.7.2 Adding and Removing Roles

You can add new role permission mappings or remove existing mappings by modifying the authorization list.

For example:

```
CREATE ROLE MY_CUSTOM_ROLE_1
GRANT PGX_SESSION_CREATE TO MY_CUSTOM_ROLE1
GRANT PGX_SERVER_GET_INFO TO MY_CUSTOM_ROLE1
GRANT MY_CUSTOM_ROLE1 TO SCOTT
```

4.2.7.3 Defining Permissions for Individual Users

In addition to defining permissions for roles, you can define permissions for individual users.

For example:

```
GRANT PGX_SESSION_CREATE TO SCOTT
GRANT PGX_SERVER_GET_INFO TO SCOTT
```

4.2.7.4 Defining Permissions to Use Custom Graph Algorithms

You can define permissions to allow developers to compile custom graph algorithms.

For example,

- Add the following static permission to the list of permissions:

```
GRANT PGX_SESSION_COMPILE_ALGORITHM TO GRAPH_DEVELOPER
```

4.2.8 Revoking Access to the Graph Server

To revoke a user's ability to access the graph server, either drop the user from the database or revoke the corresponding roles from the user, depending on how you defined the access rules in your `pgx.conf` file.

For example:

```
REVOKE graph_developer FROM scott
```

Revoking Graph Permissions

If you have the `MANAGE` permission on a graph, you can revoke graph access from users or roles using the `PgxGraph#revokePermission` API. For example:

```
PgxGraph g = ...
g.revokePermission(new PgxRole("GRAPH_DEVELOPER")) // revokes
previously granted role access
g.revokePermission(new PgxUser("SCOTT")) // revokes previously granted
user access
```

4.2.9 Examples of Custom Authorization Rules

You can define custom authorization rules for developers.

- [Example 4-1](#)
- [Example 4-2](#)
- [Example 4-3](#)

Example 4-1 Allowing Developers to Publish Graphs

Sharing of graphs with other users should be done in Oracle Database where possible. Use `GRANT` statements on the database tables so that other users can create graphs from the tables. If the graph is in the Property Graph schema use the [OPG_APIS.GRANT_ACCESS](#) API to share the graph.

In the in-memory graph server you can use the following permissions to share a graph that is already in memory, with other users connected to the graph server.

Table 4-4 Allowed Permissions

Permission	Actions Enabled by this Permission
READ	<ul style="list-style-type: none"> • READ the graph via the PGX API or in PGQL queries in PGX, create a subgraph, or clone the graph
MANAGE	<ul style="list-style-type: none"> • Publish the graph or snapshot • Includes READ and EXPORT • Grant or revoke READ and EXPORT permissions on the graph
EXPORT	<ul style="list-style-type: none"> • Export the graph to a file. • Includes READ permission.

The creator of the graph automatically gets the `MANAGE` permission granted on the graph. If you have the `MANAGE` permission, you can grant other roles or users `READ` or `EXPORT` permission on the graph. You **cannot** grant `MANAGE` on a graph. The following example of a user named `userA` shows how:

```
import oracle.pgx.api.*
import oracle.pgx.common.auth.*
```



```
PgxSession session = GraphServer.getInstance("<base-url>", "<userA>",
"<password-of-userA>").createSession("userA")
PgxGraph g = session.readGraphWithProperties("examples/sample-graph.json",
"sample-graph")
g.grantPermission(new PgxRole("GRAPH_DEVELOPER"), PgxResourcePermission.READ)
g.publish()
```

Now other users with the GRAPH_DEVELOPER role can access this graph and have READ access on it, as shown in the following example of userB:

```
PgxSession session = GraphServer.getInstance("<base-url>", "<userB>",
"<password-of-userB>").createSession("userB")
PgxGraph g = session.getGraph("sample-graph")
g.queryPsql("select count(*) from match (v)").print().close()
```

Similarly, graphs can be shared with individual users instead of roles, as shown in the following example:

```
g.grantPermission(new PgxUser("OTHER_USER"), PgxResourcePermission.EXPORT)
```

where OTHER_USER is the user name of the user that will receive the EXPORT permission on graph *g*.

Example 4-2 Allowing Developers to Access Preloaded Graphs

To allow developers to access preloaded graphs (graphs loaded during graph server startup), grant the read permission on the preloaded graph in the `pgx.conf` file. For example:

```
"preload_graphs": [{
  "path": "/data/my-graph.json",
  "name": "global_graph"
}],
"authorization": [{
  "pgx_role": "GRAPH_DEVELOPER",
  "pgx_permissions": [{
    "preloaded_graph": "global_graph"
    "grant": "read"
  }],
  ...
```

You can grant READ, EXPORT, or MANAGE permission.

Example 4-3 Allowing Developers Access to the Hadoop Distributed Filesystem (HDFS) or the Local File System

To allow developers to read files from HDFS, you must first declare the HDFS directory and then map it to a read or write permission. For example:

```
CREATE OR REPLACE DIRECTORY pgx_file_location AS 'hdfs:/data/graphs'
GRANT READ ON DIRECTORY pgx_file_location TO GRAPH_DEVELOPER
```

Similarly, you can add another permission with `GRANT WRITE` to allow write access. Such a write access is required in order to export graphs.

Access to the local file system (where the graph server runs) can be granted the same way. The only difference is that location would be an absolute file path without the `hdfs:` prefix. For example:

```
CREATE OR REPLACE DIRECTORY pgx_file_location AS '/opt/oracle/graph/
data'
```

Note that in addition to the preceding configuration, the operating system user that runs the graph server process must have the corresponding directory privileges to actually read or write into those directories.

4.2.10 Kerberos Enabled Authentication

The graph server (PGX) can authenticate users using an Oracle Database with Kerberos enabled as identity provider.

You can log into the graph server using a Kerberos ticket and the actions which you are allowed to do on the graph server are determined by the roles that have been granted to you in the Oracle Database.

- [Prerequisite Requirements](#)
- [Prepare the Graph Server for Kerberos Authentication](#)
- [Login to the Graph Server Using Kerberos Ticket](#)

4.2.10.1 Prerequisite Requirements

In order to enable Kerberos authentication on the in-memory graph server (PGX), the following system requirements must be met:

- The database needs to have Kerberos authentication enabled. See [Configuring Kerberos Authentication](#) for more information.
- Both the database and the Kerberos Authentication Server need to be reachable from the host where the graph server runs.
- The database is prepared for graph server authentication. That is, relevant graph roles have been granted to users who will log into the graph server.

4.2.10.2 Prepare the Graph Server for Kerberos Authentication

The following are the steps to enable Kerberos authentication on the in-memory graph server (PGX):

1. Locate the `pgx.conf` file of your installation.

 **Note:**

If you installed the graph server via RPM, the file is located at: `/etc/oracle/graph/pgx.conf`

2. Locate the `krb5_conf_file` line of the realm options, inside the `pgx.conf` file:

```
"pgx_realm": {
  "implementation": "oracle.pg.identity.DatabaseRealm",
  "options": {
    ...
    "krb5_conf_file": "<REPLACE-WITH-KRB5-CONF-FILE-PATH-TO-ENABLE-
KERBEROS-AUTHENTICATION>",
    "krb5_ticket_cache_dir": "/dev/shm",
    "krb5_max_cache_size": 1024
  }
},
```

3. Replace the text with the `krb5.conf` file that you are using for the database and user authentication. For example:

```
"pgx_realm": {
  "implementation": "oracle.pg.identity.DatabaseRealm",
  "options": {
    ...
    "krb5_conf_file": "/etc/krb5.conf",
    "krb5_ticket_cache_dir": "/dev/shm",
    "krb5_max_cache_size": 1024
  }
},
```

 **Note:**

The file provided for the `krb5_conf_file` option needs to be valid and readable by the graph server. In case you don't replace the `krb5_conf_file` value or the value is empty, then the graph server will not use Kerberos authentication.

Also, you can set the cache directory that will be used for the graph server to temporarily store Kerberos tickets given by clients as well as the maximum cache size after which new login attempts will be rejected. The cache size represents the maximum amount of concurrent Kerberos sessions active on the graph server.

4.2.10.3 Login to the Graph Server Using Kerberos Ticket

The following are the steps to login to the in-memory graph server (PGX) using Kerberos ticket:

1. Create a new Kerberos ticket using the `okinit` command:

```
$ okinit <username>
```

This will prompt for your password and then create a new Kerberos ticket.

2. Connect to a remote graph server with only the base URL parameter using JShell:

```
$ opg4j -b https://localhost:7007
```

Or using Python client:

```
$ opg4py -b https://localhost:7007
```

On Linux, JShell and Python interactive client shells automatically detect the Kerberos ticket on your local file system and use that to authenticate with the graph server.

3. In case the auto-detection is not working, you can also explicitly pass in the ticket to the shell. Run the `oklist` command, to find the location of the ticket on the local file system.

```
$ oklist
```

```
Kerberos Utilities for Linux: Version 19.0.0.0.0 - Production on 31-  
MAR-2021 15:26:46
```

```
Copyright (c) 1996, 2019 Oracle. All rights reserved.
```

```
Configuration file : /etc/krb5.conf.  
Ticket cache: FILE:/tmp/krb5cc_54321  
Default principal: oracle@realm
```

4. Specify your Kerberos ticket path using the `--kerberos_ticket` parameter. For example, using JShell:

```
$ opg4j -b https://localhost:7007 --kerberos_ticket /tmp/  
krb5cc_54321
```

Or using Python Client:

```
$ opg4py -b https://localhost:7007 --kerberos_ticket /tmp/  
krb5cc_54321
```

If you are using a Java client program (or JShell on embedded mode), you can get a server instance using the following API:

```
...  
ServerInstance instance = GraphServer.getInstance("https://  
localhost:7007", "/tmp/krb5cc_54321");  
PgxSession session = instance.createSession("my-session");  
...
```

If you are using a Python Client program (or `opg4py` on embedded mode), you can get a server instance using the following API

```
...  
instance = graph_server.get_instance("https://localhost:7007",  
"/tmp/krb5cc_54321")  
session = instance.create_session("my-session")  
...
```

If you are connecting to a remote graph server, all you need is the Oracle Graph Client to be installed. For example:

```
import sys
import pypgx as pgx

sys.path.append("/path/to/graph/client/oracle-graph-client-21.2.0/python/
pypgx/pg/rdbms")

import graph_server

base_url = "https://localhost:7007"
kerberos_ticket = "/tmp/krb5cc_54321"

instance = graph_server.get_instance(base_url, kerberos_ticket)
print(instance)
```

4.3 About Vertex and Edge IDs

Generating vertex and edge IDs when loading from database tables into PGX

PGX enforces by default the existence of a unique identifier for each vertex and edge in a graph, so that they can be retrieved by using `PgxGraph.getVertex(ID id)` and `PgxGraph.getEdge(ID id)` or by PGQL queries using the built-in `id()` method.

The following supported ID generation strategies can be selected through the configuration parameters `vertex_id_strategy` and `edge_id_strategy`:

- `keys_as_ids`: This is the default strategy to generate vertex IDs.
- `partitioned_ids`: This is the recommended strategy for partitioned graphs.
- `unstable_generated_ids`: This results in system generated vertex or edge IDs.
- `no_ids`: This strategy disables vertex or edge IDs and therefore prevents you from calling APIs using vertex or edge IDs.

Using keys to generate IDs

The default strategy to generate the vertex IDs is to use the keys provided during loading of the graph (`keys_as_ids`). In that case, each vertex should have a vertex key that is unique across all providers.

For edges, by default no keys are required in the edge data, and edge IDs will be automatically generated by PGX (`unstable_generated_ids`). This automatic ID generation can be applied for vertex IDs also. Note that the generation of vertex or edge IDs is not guaranteed to be deterministic. If required, it is also possible to load edge keys as IDs.

The `partitioned_ids` strategy requires keys to be unique only *within* a vertex or edge provider (data source). The keys do not have to be globally unique. Globally unique IDs are derived from a combination of the provider name and the key inside the provider, as `<provider_name><unique_key_within_provider>`. For example, `Account(1)`.

The `partitioned_ids` strategy can be set through the configuration fields `vertex_id_strategy` and `edge_id_strategy`. For example,

```
{
  "name": "bank_graph_analytics",
  "optimized_for": "updates",
  "vertex_id_strategy" : "partitioned_ids",
  "edge_id_strategy" : "partitioned_ids",
  "vertex_providers": [
    {
      "name": "Accounts",
      "format": "rdbms",
      "database_table_name": "BANK_ACCOUNTS",
      "key_column": "ID",
      "key_type": "integer",
      "props": [
        {
          "name": "ID",
          "type": "integer"
        },
        {
          "name": "NAME",
          "type": "string"
        }
      ],
      "loading": {
        "create_key_mapping" : true
      }
    }
  ],
  "edge_providers": [
    {
      "name": "Transfers",
      "format": "rdbms",
      "database_table_name": "BANK_TXNS",
      "key_column": "ID",
      "source_column": "FROM_ACCT_ID",
      "destination_column": "TO_ACCT_ID",
      "source_vertex_provider": "Accounts",
      "destination_vertex_provider": "Accounts",
      "props": [
        {
          "name": "ID",
          "type": "integer"
        },
        {
          "name": "AMOUNT",
          "type": "double"
        }
      ],
      "loading": {
        "create_key_mapping" : true
      }
    }
  ]
}
```

```
]
}
```

**Note:**

All available key types are supported in combination with partitioned IDs.

After the graph is loaded, PGX maintains information about which property of a provider corresponds to the key of the provider. In the preceding example, the vertex property `ID` happens to correspond to the vertex key and also the edge property `ID` happens to correspond to the edge key. Each provider can have at most one such "key property" and the property can have any name.

Key properties are used for internal optimizations as well as for providing keys for the vertex or edge or both when inserting new entities. Key properties are currently non-updatable. Trying to update a key property will result in an error. For example,

```
vertex key property ID cannot be updated
```

Using an auto-incrementer to generate partitioned IDs

It is recommended to always set `create_key_mapping` to `true` to benefit from performance optimizations. But if there are no single-column keys for edges, `create_key_mapping` can be set to `false`. Similarly, `create_key_mapping` can be set to `false` for vertex providers also. IDs will be generated via an auto-incrementer, for example `Accounts(1)`, `Accounts(2)`, `Accounts(3)`.

See [PGQL Queries with Partitioned IDs](#) for more information on executing PGQL queries with partitioned IDs.

4.4 Reading Graphs from Oracle Database into the Graph Server (PGX)

Once logged in to the graph server (PGX), you can now read graphs from the database into the graph server without specifying any connection information in the graph configuration.

Your database user must exist and have read access on the graph data in the database.

There are several ways to read a graph into the graph server (PGX) from Oracle Database:

- Using a Property Graph View
See [Loading a Graph into the Graph Server \(PGX\) from a Property Graph View](#) for more information.
- Using the PGQL - `CREATE PROPERTY GRAPH` statement.
See [Creating a Property Graph using PGQL](#) for more details.
- Using a PGX graph configuration file in JSON format
See [Creating a JSON Configuration to Load a Graph](#) for more details.
- Using [GraphConfigBuilder](#) class to create Oracle RDBMS graph configs programmatically via Java methods.
See [Building Graph Configuration Using the GraphConfigBuilder API](#) for more details.

- [Creating a JSON Configuration to Load a Graph](#)
- [Building Graph Configuration Using the GraphConfigBuilder API](#)

4.4.1 Creating a JSON Configuration to Load a Graph

Loading a Graph from the Property Graph Schema Using JSON Configuration

The graph server (PGX) allows the loading of property graphs from Oracle Database using the property graph schema.

You can create a graph configuration file to load the graph. The following shows a sample graph configuration file that uses a keystore to connect to the database. See [Store the Database Password in a Keystore](#) for more information on using a keystore.

```
{
  "name": "BANK_ACCOUNTS",
  "format": "pg",
  "db_engine": "rdbms",
  "jdbc_url": "jdbc:oracle:thin:@localhost:1521/orclpdb",
  "username": "graphuser",
  "keystore_alias": "database1",
  "vertex_props": [
    {
      "name": "ID",
      "type": "integer"
    },
    {
      "name": "NAME",
      "type": "string"
    }
  ],
  "edge_props": [
    {
      "name": "AMOUNT",
      "type": "float"
    },
    {
      "name": "DESCRIPTION",
      "type": "string"
    }
  ],
  "partition_while_loading": "by_label",
  "loading": {
    "load_edge_label": true,
    "load_vertex_labels": true,
    "create_vertex_id_mapping": true,
    "create_edge_id_mapping": true
  }
}
```

See [API for Loading Graphs into Memory](#) for more information on `PgxSession` API methods for reading graphs into memory.

You can now read the graph into the graph server (PGX) using the `PgxSession` API method as shown:

-
- [JShell](#)
 - [Java](#)
 - [Python](#)

JShell

```
opg4j> var g = session.readGraphWithProperties("bank_graph.json")
```

Java

```
PgxGraph g = session.readGraphWithProperties("bank_graph.json");
```

Python

```
g = session.read_graph_with_properties("bank_graph.json")
```

Loading a Partitioned Graph Using JSON Configuration

In order to load a partitioned graph into the graph server (PGX), you can create a graph configuration file, which contains metadata about the graph to be loaded.

The following shows a sample partitioned graph configuration file:

```
{
  "name": "bank_graph_analytics",
  "vertex_providers": [
    {
      "name": "Accounts",
      "format": "rdbms",
      "database_table_name": "BANK_NODES",
      "key_column": "ID",
      "key_type": "integer"
    }
  ],
  "edge_providers": [
    {
      "name": "Transfers",
      "format": "rdbms",
      "database_table_name": "BANK_EDGES_AMT",
      "key_column": "ID",
      "source_column": "SRC_ID",
      "destination_column": "DEST_ID",
      "source_vertex_provider": "Accounts",
      "destination_vertex_provider": "Accounts",
    }
  ]
}
```

```
        "props": [
            {
                "name": "AMOUNT",
                "type": "float"
            }
        ]
    }
]
```

See [Graph Configuration Options](#) for more details on partitioned graph configuration options.

You can now read the graph into the graph server as shown:

-
- [JShell](#)
 - [Java](#)
 - [Python](#)

JShell

```
opg4j> var g =
session.readGraphWithProperties("bank_graph_analytics.json")
```

Java

```
PgxGraph g =
session.readGraphWithProperties("bank_graph_analytics.json");
```

Python

```
g = session.read_graph_with_properties("bank_graph_analytics.json")
```

4.4.2 Building Graph Configuration Using the GraphConfigBuilder API

You can load a graph from Oracle Database by first defining the graph configuration object using the [GraphConfigBuilder](#) class and then reading the graph into the graph server (PGX).

The following are a few examples to read a property graph into memory, authenticating as <database user>/<database password> with the database:

Loading a Graph from Property Graph Schema Using `GraphConfigBuilder`

- JShell
- Java

JShell

```
opg4j> var config = GraphConfigBuilder.forPropertyGraphRdbms().
...>     setName("bank_graph").
...>     setJdbcUrl("jdbc:oracle:thin:@localhost:1521/orclpdb").
...>     setUsername("graphuser").
...>     setPassword("graphuser").
...>     addVertexProperty("ID", PropertyType.STRING).
...>     addVertexProperty("NAME", PropertyType.STRING).
...>     addEdgeProperty("AMOUNT", PropertyType.FLOAT).
...>     addEdgeProperty("DESCRIPTION", PropertyType.STRING).
...>     setPartitionWhileLoading(PartitionWhileLoading.BY_LABEL).
...>     setLoadVertexLabels(true).
...>     setLoadEdgeLabel(true).
...>     build()
config ==> {"error_handling":
{ }, "name": "bank_graph", "vertex_id_type": "long", "loading":
{"load_vertex_labels": true, "load_edge_label": true}, "password": "*****", "form
at": "pg", "edge_props": [{"name": "AMOUNT", "dimension": 0, "type": "float"},
{"name": "DESCRIPTION", "dimension": 0, "type": "string"}], "jdbc_url": "jdbc:oracle
:thin:@localhost:1521/orclpdb", "vertex_props":
[{"name": "ID", "dimension": 0, "type": "string"},
{"name": "NAME", "dimension": 0, "type": "string"}], "partition_while_loading": "BY_
LABEL", "attributes": { }, "db_engine": "RDBMS", "username": "graphuser"}
opg4j> PgxGraph g = session.readGraphWithProperties(config)
g ==> PgxGraph[name=bank_graph,N=1000,E=5001,created=1632825497751]
```

Java

```
GraphConfig config = GraphConfigBuilder.forPropertyGraphRdbms()
    .setName("bank_graph")
    .setJdbcUrl("jdbc:oracle:thin:@localhost:152
1/orclpdb")
    .setUsername("graphuser")
    .setPassword("graphuser")
    .addVertexProperty("ID",
PropertyType.STRING)
    .addVertexProperty("NAME",
PropertyType.STRING)
    .addEdgeProperty("AMOUNT",
PropertyType.FLOAT)
    .addEdgeProperty("DESCRIPTION",
PropertyType.STRING)
    .setPartitionWhileLoading(PartitionWhileLoad
ing.BY_LABEL)
    .setLoadVertexLabels(true)
    .setLoadEdgeLabel(true)
    .build();
```

```
PgxGraph g = session.readGraphWithProperties(config);
```

Loading a Partitioned Graph into the Graph Server (PGX) Using GraphConfigBuilder

- [JShell](#)
- [Java](#)

JShell

```
opg4j> var vertexConfig = new RdbmsEntityProviderConfigBuilder(). //
build the vertex provider
...>                                     setName("node").
...>                                     setKeyColumn("ID").
...>
setDatabaseTableName("BANK_ACCOUNTS").
...>                                     addProperty("ID",
PropertyType.INTEGER).
...>                                     build()
vertexConfig ==> {"name":"node","error_handling":
{"key_type":"long","attributes":{},"props":
[{"dimension":0,"name":"ID","type":"integer"}],"loading":
{"create_key_mapping":true},"format":"rdbms",
"database_table_name":"BANK_ACCOUNTS","key_column":"ID"}
opg4j> var edgeConfig = new RdbmsEntityProviderConfigBuilder(). //
build the edge provider
...>                                     setName("transfer").
...>                                     setKeyColumn("ID").
...>
setSourceColumn("FROM_ACCT_ID").
...>
setDestinationColumn("TO_ACCT_ID").
...>
setSourceVertexProvider("node").
...>
setDestinationVertexProvider("node").
...>
createKeyMapping(true).
...>
setDatabaseTableName("BANK_TXNS").
...>
addProperty("FROM_ACCT_ID", PropertyType.LONG).
...>
addProperty("TO_ACCT_ID", PropertyType.LONG).
...>
addProperty("AMOUNT", PropertyType.LONG).
...>                                     build()
edgeConfig ==>
```

```

{"name":"transfer","destination_vertex_provider":"node","attributes":
 {}, "props":[{"dimension":0,"name":"FROM_ACCT_ID","type":"long"},
 {"dimension":0,"name":"TO_ACCT_ID","type":"long"},
 {"dimension":0,"name":"AMOUNT","type":"long"}],"loading":
 {"create_key_mapping":true,"database_table_name":"BANK_TXNS","key_column":"I
 D","destination_column":"TO_ACCT_ID","source_vertex_provider":"node",
 "error_handling":
 {}, "key_type":"long","source_column":"FROM_ACCT_ID","format":"rdbms"}
opg4j> var cfg = GraphConfigBuilder.forPartitioned(). //build the
partitioned graph
...>                                setJdbcUrl("jdbc:oracle:thin:@localhost:1521/
orclpdb").
...>                                setUsername("graphuser").
...>                                setPassword("graphuser").
...>                                setName("bank_graph_partitioned").
...>                                setVertexIdType(IdType.LONG).
...>                                addVertexProvider(vertexConfig).
...>                                addEdgeProvider(edgeConfig).
...>                                build()

opg4j> var g = session.readGraphWithProperties(cfg)
g ==>
PgxGraph[name=bank_graph_partitioned_2,N=1000,E=5001,created=1632837214584]

```

Java

```

// Build the vertex provider
RdbmsEntityProviderConfig vertexConfig = new
RdbmsEntityProviderConfigBuilder()
                                .setName("node")
                                .setKeyColumn("ID")
                                .setDatabaseTableName("BANK_ACCO
UNTS")
                                .addProperty("ID",
PropertyType.INTEGER)
                                .build();

// Build the edge provider
RdbmsEntityProviderConfig edgeConfig = new RdbmsEntityProviderConfigBuilder()
                                .setName("transfer")
                                .setKeyColumn("ID")
                                .setSourceColumn("FROM_ACCT_ID")
                                .setDestinationColumn("TO_ACCT_I
D")
                                .setSourceVertexProvider("node")
                                .setDestinationVertexProvider("n
ode")
                                .createKeyMapping(true)
                                .setDatabaseTableName("BANK_TXNS
")
                                .addProperty("FROM_ACCT_ID",
PropertyType.LONG)
                                .addProperty("TO_ACCT_ID",
PropertyType.LONG)
                                .addProperty("AMOUNT",

```

```

PropertyType.LONG)
                                                    .build();

// Build the partitioned graph
GraphConfig cfg = GraphConfigBuilder.forPartitioned()
    .setJdbcUrl("jdbc:oracle:thin:@<host>:<port>/
<service_name>")
    .setUsername("<username>")
    .setPassword("<password>")
    .setName("bank_graph_partitioned")
    .setVertexIdType(IdType.LONG)
    .addVertexProvider(vertexConfig)
    .addEdgeProvider(edgeConfig)
    .build();

PgxGraph g = session.readGraphWithProperties(cfg);

```

4.5 Keeping the Graph in Oracle Database Synchronized with the Graph Server

You can use the `FlashbackSynchronizer` API to automatically apply changes made to graph in the database to the corresponding `PgxGraph` object in memory, thus keeping both synchronized.

This API uses Oracle's Flashback Technology to fetch the changes in the database since the last fetch and then push those changes into the graph server using the `ChangeSet` API. After the changes are applied, the usual snapshot semantics of the graph server apply: each delta fetch application creates a new in-memory snapshot. Any queries or algorithms that are executing concurrently to snapshot creation are unaffected by the changes until the corresponding session refreshes its `PgxGraph` object to the latest state by calling the `session.setSnapshot(graph, PgxCSession.LATEST_SNAPSHOT)` procedure.

For detailed information about Oracle Flashback technology, see the Database Development Guide.

Prerequisites for Synchronizing

The Oracle database must have Flashback enabled and the database user that you use to perform synchronization must have:

- Read access to all tables which need to be kept synchronized.
- Permission to use flashback APIs. For example:

```
GRANT EXECUTE ON DBMS_FLASHBACK TO <user>
```

The database must also be configured to retain changes for the amount of time needed by your use case.

Types of graphs that can be synchronized

Not all `PgxGraph` objects in PGX can be synchronized. The following limitations apply:

- Only the original creator of the graph can synchronize it. That is, the current user must have the `MANAGE` permission of the graph.
- Only graphs loaded from database tables ("partitioned graphs") can be synchronized. Graphs created from other formats or graphs created via the graph builder API cannot be synchronized.
- Only the *latest snapshot* of a graph can be synchronized.

Types of changes that can be synchronized

The synchronizer supports keeping the in-memory graph snapshot in sync with the following database-side modifications:

- insertion of new vertices and edges
- removal of existing vertices and edges
- update of property values of any vertex or edge

The synchronizer does not support schema-level changes to the input graph, such as:

- alteration of the list of input vertex or edge tables
- alteration of any columns of any input tables (vertex or edge tables)

Furthermore, the synchronizer does not support updates to vertex and edge keys.

For detailed examples, see the following topic:

- [Examples of Synchronizing](#)

4.5.1 Examples of Synchronizing

You can perform your graph synchronization using the following examples:

- [Example 4-4](#)
- [Example 4-5](#)

Example 4-4 Synchronizing Graphs Using `CREATE PROPERTY GRAPH` Statement

1. Assume you have the following Oracle Database tables, `PERSONS` and `FRIENDSHIPS`.

```
CREATE TABLE persons (
    person_id NUMBER GENERATED ALWAYS AS IDENTITY (START WITH 1 INCREMENT
BY 1),
    name VARCHAR2(200),
    birthdate DATE,
    height FLOAT DEFAULT on null 0,
    CONSTRAINT person_pk PRIMARY KEY (person_id)
);
```

```
CREATE TABLE friendships (
    friendship_id NUMBER GENERATED ALWAYS AS IDENTITY (START WITH 1
INCREMENT BY 1),
    person_a NUMBER,
    person_b NUMBER,
    meeting_date DATE,
    CONSTRAINT fk_person_a_id FOREIGN KEY (person_a) REFERENCES
persons(person_id),
```

```

        CONSTRAINT fk_person_b_id FOREIGN KEY (person_b) REFERENCES
        persons(person_id)
        CONSTRAINT fs_pk PRIMARY KEY (friendship_id)
    );

```

2. You can add some sample data into these tables as shown:

```

INSERT INTO persons (name, height, birthdate) VALUES ('John', 1.80,
to_date('13/06/1963', 'DD/MM/YYYY'));
INSERT INTO persons (name, height, birthdate) VALUES ('Mary', 1.65,
to_date('25/09/1982', 'DD/MM/YYYY'));
INSERT INTO persons (name, height, birthdate) VALUES ('Bob', 1.75,
to_date('11/03/1966', 'DD/MM/YYYY'));
INSERT INTO persons (name, height, birthdate) VALUES ('Alice',
1.70, to_date('01/02/1987', 'DD/MM/YYYY'));

```

```

INSERT INTO friendships (person_a, person_b, meeting_date) VALUES
(1, 3, to_date('01/09/1972', 'DD/MM/YYYY'));
INSERT INTO friendships (person_a, person_b, meeting_date) VALUES
(2, 4, to_date('19/09/1992', 'DD/MM/YYYY'));
INSERT INTO friendships (person_a, person_b, meeting_date) VALUES
(4, 2, to_date('19/09/1992', 'DD/MM/YYYY'));
INSERT INTO friendships (person_a, person_b, meeting_date) VALUES
(3, 2, to_date('10/07/2001', 'DD/MM/YYYY'));

```

3. Write the corresponding CREATE PROPERTY GRAPH statement which describes how to load those tables as a graph as shown in the following Java code example:

```

session.executePgql(
    "CREATE PROPERTY GRAPH friends VERTEX TABLES ("
        + "  persons KEY (person_id) LABEL person PROPERTIES"
        + "  (name,height,birthdate)"
        + " )"
    + "EDGE TABLES ("
        + "  friendships "
        + "    KEY (friendship_id) "
        + "    SOURCE KEY (person_a) REFERENCES persons "
        + "    DESTINATION KEY (person_b) REFERENCES persons "
        + "    LABEL friendof PROPERTIES (meeting_date)"
        + " )"
);
PgxGraph graph = session.getGraph("friends");

```

This creates a snapshot of the graph which is loaded into memory. You can now run algorithms and queries on the graph.

4. Now change the data in the input tables in the database. For example, add new persons to the PERSONS table and also add another edge.

You can open a new JDBC connection to the database and run a few INSERT statements as shown in the following code:

```

Connection conn = DriverManager.getConnection("<jdbc-url>",
"<user>", "<pass>");
conn.createStatement().executeQuery("INSERT INTO persons (name,

```



```

birthdate, height) VALUES ('Mariana',to_date('21/08/1996','DD/MM/
YYYY'),1.65)");
conn.createStatement().executeQuery("INSERT INTO persons (name,
birthdate, height) VALUES ('Francisco',to_date('13/06/1963','DD/MM/
YYYY'),1.75)");
conn.createStatement().executeQuery("INSERT INTO friendships (person_a,
person_b, meeting_date) VALUES (1, 6, to_date('13/06/2013','DD/MM/
YYYY'))");
conn.commit();

```

Committing the changes to the database causes the graph in memory to become out of sync with the database source tables.

5. You can synchronize the in-memory graph with the database by creating a new synchronizer object as shown in the following code:

```

Synchronizer synchronizer = new
Synchronizer.Builder<FlashbackSynchronizer>()
    .setType(FlashbackSynchronizer.class)
    .setGraph(graph)
    .setConnection(conn)
    .build();

```

Internally, the graph server keeps track of the Oracle system change number (SCN) the current graph snapshot belongs to. The synchronizer is a *client-side* component which connects to the database, detects changes by comparing state of the the original input tables using the current SCN via the flashback mechanism and then sends any changes to the graph server using the changeset API. In order to do so, the synchronizer needs to know how to connect to the database (`conn` parameter) as well as which graph to keep in sync (`graph` parameter).

- Alternatively, you can use this equivalent shortcut:

```

Synchronizer synchronizer =
graph.createSynchronizer(FlashbackSynchronizer.class, conn);

```

6. Call the `sync()` operation, to fetch the database changes and create a new in-memory snapshot:

```
graph = synchronizer.sync();
```

You will notice that the two new vertices and the new edge have been applied to the graph:

```
graph ==> PgxGraph[name=FRIENDS,N=6,E=5,created=1594754376861]
```

Splitting the Fetching and Applying of Changes

The `synchronizer.sync()` invocation in the preceding code, fetches the changes and applies them in one call. However, you can encode a more complex update logic by splitting this process into separate `fetch()` and `apply()` invocations. For example:

```
synchronizer.fetch() // fetches changes from the database
if (synchronizer.getGraphDelta().getTotalNumberOfChanges() > 100)
{ // only create snapshot if there have been more than 100 changes
  synchronizer.apply()
}
```

Example 4-5 Synchronizing Graphs Created Via Graph Configuration Objects

[Example 4-4](#) uses a `CREATE PROPERTY GRAPH` statement to create the graph which hides some of the more advanced graph configuration options.

Though synchronization of graphs created via graph configuration objects is supported in general, the following few limitations apply:

- Only partitioned graph configurations with all providers being database tables are supported.
- Each edge or vertex provider or both must specify the owner of the table by setting the `username` field. For example, if user `SCOTT` owns the table, then set the `username` accordingly in the provider block of that table:

```
"username": "scott"
```

- In the root loading block, the snapshot source must be set to `change_set`:

```
"loading": {
  "snapshots_source": "change_set"
}
```

- It is highly recommended to set the `optimized_for` field to `updates` to avoid memory exhaustion when creating many snapshots:

```
"optimized_for": "updates"
```

You can load the same graph shown in [Example 4-4](#) using the following graph configuration (JSON) file:

```
{
  "name": "friends",
  "optimized_for": "updates",
  "vertex_id_strategy": "partitioned_ids",
  "edge_id_strategy": "partitioned_ids",
  "edge_id_type": "long",
  "vertex_id_type": "long",
  "jdbc_url": "<jdbc_url>",
  "username": "<username>",
  "keystore_alias": "<keystore_alias>",
  "vertex_providers": [
    {
      "format": "rdbms",
```

```

    "username": "<username>",
    "key_type": "long",
    "name": "person",
    "database_table_name": "persons",
    "key_column": "person_id",
    "props": [
        ...
    ],
    "loading": {
        "create_key_mapping": true
    }
},
"edge_providers": [
    {
        "format": "rdbms",
        "username": "<username>",
        "name": "friendOf",
        "source_vertex_provider": "person",
        "destination_vertex_provider": "person",
        "database_table_name": "friendships",
        "source_column": "person_a",
        "destination_column": "person_b",
        "key_column": "friendship_id",
        "key_type": "long",
        "props": [
            ...
        ],
        "loading": {
            "create_key_mapping": true
        }
    }
],
"loading": {
    "snapshots_source": "change_set"
}
}

```

 **Note:**

- In the preceding JSON file, replace the values <jdbc_url>, <username>, and <keystore_alias> with the values for connecting to your database.
- When using the graph configuration file, you can load the graph into memory using JShell (be sure to register the keystore containing the database password when starting it) :

```

var pgxGraph =
session.readGraphWithProperties("<name_of_config_file>.json");

```

4.6 Optimizing Graphs for Read Versus Updates in the In-Memory Graph Server (PGX)

The in-memory graph server (PGX) can store an optimized graph for other reads or updates. This is only relevant when the updates are made directly to a graph instance in the in-memory graph server.

Graph Optimized for Reads

Graphs optimized for reads will provide the best performance for graph analytics and PGQL queries. In this case there could be potentially higher latencies to update the graph (adding or removing vertex and edges or updating the property values of previously existing vertex or edges through `GraphChangeSet` API). There could also be higher memory consumption. When using graphs optimized for reads, each updated graph or graph snapshot consumes memory proportional to the size of the graph in terms of vertices and edges.

The `optimized_for` configuration property can be set to `reads` when loading the graph into the in-memory graph server (PGX) to create a graph instance that is optimized for reads.

Graph Optimized for Updates

Graphs optimized for updates use a representation enabling low-latency update of graphs. With this representation, the graph server can reach millisecond-scale latencies when updating graphs with millions of vertices and edges (this is indicative and will vary depending on the hardware configuration).

To achieve faster update operations, graph server avoids as much as possible doing a full duplication of the previous graph (snapshot) to create a new graph (snapshot). This also improves the memory consumption (in typical scenarios). New snapshots (or new graphs) will only consume additional memory proportional to the memory required for the changes applied.

In this representation, there could be lower performance of graph queries and analytics.

The `optimized_for` configuration property can be set to `updates` when loading the graph into the in-memory graph server (PGX) to create a graph instance that is optimized for reads.

4.7 Storing a Graph Snapshot on Disk

After reading a graph into memory using either Java or the Shell, if you make some changes to the graph such as running the PageRank algorithm and storing the values as vertex properties, you can store this snapshot of the graph on disk.

This is helpful if you want to save the state of the graph in memory, such as if you must shut down the in-memory graph server to migrate to a newer version, or if you must shut it down for some other reason.

(Storing graphs over HTTP/REST is currently not supported.)

A snapshot of a graph can be saved as a file in a binary format (called a PGB file) if you want to save the state of the graph in memory, such as if you must shut down the in-memory graph server to migrate to a newer version, or if you must shut it down for some other reason.

In general, we recommend that you store the graph queries and analytics APIs that had been executed, and that after the in-memory graph server has been restarted, you reload and re-execute the APIs. But if you must save the state of the graph, you can use the logic in the following example to save the graph snapshot from the shell.

In a three-tier deployment, the file is written on the server-side file system. You must also ensure that the file location to write is specified in the in-memory graph server. (As explained in [Three-Tier Deployments of Oracle Graph with Autonomous Database](#), in a three-tier deployment, access to the PGX server file system requires a list of allowed locations to be specified.)

```

opg4j> var graph =
session.createGraphBuilder().addVertex(1).addVertex(2).addVertex(3).addEdge(1
,2).addEdge(2,3).addEdge(3, 1).build()
graph ==> PgxGraph[name=anonymous_graph_1,N=3,E=3,created=1581623669674]

opg4j> analyst.pagerank(graph)
$3 ==> VertexProperty[name=pagerank,type=double,graph=anonymous_graph_1]

// Now save the state of this graph
opg4j> graph.store(Format.PGB, "/scratch/PG/snapshot/snapshot.pgb")
$16 ==> {"edge_uris":[],"error_handling":{},"attributes":
{},"format":"pgb","vertex_props":
[{"type":"double","name":"pagerank","dimension":0},"edge_props":[],"vertex_
id_type":"integer","vertex_uris":["/scratch/PG/snapshot/
snapshot.pgb"],"loading":{}}

// reload from disk
opg4j> var graphFromDisk = session.readGraphFile("/scratch/PG/snapshot/
snapshot.pgb", Format.PGB)
graphFromDisk ==> PgxGraph[name=snapshot,N=3,E=3,created=1635791377443]

// previously computed properties are still part of the graph and can be
queried
opg4j> graphFromDisk.queryPgql("SELECT x.pagerank MATCH (x)").print().close()
+-----+
| x.pagerank          |
+-----+
| 0.3333333333333333 |
| 0.3333333333333333 |
| 0.3333333333333333 |
+-----+

```

The following example is essentially the same as the preceding one, but it uses partitioned graphs. You must first load a partitioned graph into the graph server (PGX). See [Loading a Partitioned Graph into the Graph Server \(PGX\) Using GraphConfigBuilder](#) for an example.

```

opg4j> analyst.pagerank(graph)
$8 ==> VertexProperty[name=pagerank,type=double,graph=bank_graph_partitioned]

// Now save the state of this graph

```

```

opg4j> var storedPgbConfig = graph.store(ProviderFormat.PGB, "/
scratch/PG/snapshot/snapshot_")
storedPgbConfig ==> {"vertex_id_strategy":"KEYS_AS_IDS","loading":
{},"array_compaction_threshold":0.2,"time_with_timezone_format":
["h[h]:m[m]:s[s] a[ XXX]","[
yyyy-MM-dd'T']H[H]:m[m]:s[s][.SSSSSS]][XXX]","yyyy-MM-dd'T'H[H]:m[m]
[:s[s][.SSSSSS]]
[XXX]"],"name":"bank_graph_partitioned","timestamp_with_timezone_format
":
["yyyy-MM-dd'T'H[H]:m[m]:s[s][.SSSSSS]][XXX]","yyyy-MM-dd H[H]:m[m]
[:s[s][.SSSSSS]][XXX]","yyyy-MM-dd'T'H[H]:m[m]:s[s][.SSSSSS]]
[XXX]"],"edge_providers":{
{"loading":
{"create_key_mapping":true},"destination_vertex_provider":"node","desti
nation_column":3,"format":"pgb","time_format":["h[h]:m[m]:s[s][.SSS]
a[ XXX]
","[yyyy-MM-dd'T']H[H]:m[m]:s[s][.SSSSSS]][XXX]","yyyy-MM-
dd'T'H[H]:m[m]:s[s][.SSSSSS]]
[XXX]"],"key_type":"long","has_keys":true,"label":"transfer","source_
column":2,"storing":{"base_path":"/scratch/PG/snapshot/
snapshot_transfer","edge_extension":"pgb"},"attributes":
{},"timestamp_format":["yyyy-MM-dd'T'H[H]:m[m]:s
[s][.SSSSSS]][XXX]","yyyy-MM-dd H[H]:m[m]:s[s][.SSSSSS]]
[XXX]","yyyy-MM-dd'T'H[H]:m[m]:s[s][.SSSSSS]]
[XXX]"],"time_with_timezone_format":["h[h]:m[m]:s[s]
] a[ XXX]","[yyyy-MM-dd'T']H[H]:m[m]:s[s][.SSSSSS]][XXX]","yyyy-MM-
dd'T'H[H]:m[m]:s[s][.SSSSSS]][XXX]"],"timestamp_with_timezone_format":
["yyyy-MM-dd'T'H[H]
:m[m]:s[s][.SSSSSS]][XXX]","yyyy-MM-dd H[H]:m[m]:s[s][.SSSSSS]]
[XXX]","yyyy-MM-dd'T'H[H]:m[m]:s[s][.SSSSSS]][XXX]"],"props":
[{"drop_after_loading":false,
"type":"long","column":4,"name":"FROM_ACCT_ID","dimension":0},
{"drop_after_loading":false,"type":"long","column":5,"name":"TO_ACCT_ID
","dimension":0},{"drop_aft
er_loading":false,"type":"long","column":6,"name":"AMOUNT","dimension":
0}], "header":false,"source_vertex_provider":"node","error_handling":
{},"uris":["/scratch/
PG/snapshot/snapshot_transfer.pgb"],"key_column":1,"local_date_format":
["yyyy-M[M]-d[d]","M[M]/d[d]/yyyy","d[d]-MMM-yyyy","d[d]-M[M]-
yyyy","yyyy-MM-dd'T'H[H]:
m[m]:s[s][.SSSSSS]][XXX]","yyyy-MM-dd'T'H[H]:m[m]:s[s][.SSSSSS]]
[XXX]"],"name":"transfer"},"local_date_format":["yyyy-M[M]-
d[d]","M[M]/d[d]/yyyy","d[d]-MMM
-yyyy","d[d]-M[M]-yyyy","yyyy-MM-dd'T'H[H]:m[m]:s[s][.SSSSSS]]
[XXX]","yyyy-MM-dd'T'H[H]:m[m]:s[s][.SSSSSS]]
[XXX]"],"vertex_id_type":"long","optimized_for":
"READ","max_prefetched_rows":10000,"num_connections":2,"max_batch_size"
:10000,"error_handling":
{"on_missing_vertex":"ERROR"},"scroll_time":"1m","time_format":["
h[h]:m[m]:s[s][.SSS] a[ XXX]","[yyyy-MM-dd'T']H[H]:m[m]:s[s]
[.SSSSSS]][XXX]","yyyy-MM-dd'T'H[H]:m[m]:s[s][.SSSSSS]]
[XXX]"],"edge_id_strategy":"KEYS_AS_ID
S","attributes":{},"vertex_providers":[{"storing":
{"vertex_extension":"pgb","base_path":"/scratch/PG/snapshot/
snapshot_node"},"attributes":{},"loading":{"create

```

```

_key_mapping":true}, "timestamp_format":["yyyy-MM-dd'T'H[H]:m[m][:s[s]
[.SSS[SSS]]][XXX]", "yyyy-MM-dd H[H]:m[m][:s[s][.SSS[SSS]]][XXX]", "yyyy-MM-
dd'T'H[H]:m[m][:s
[s][.SSSSSS][XXX]"], "time_with_timezone_format":["h[h]:m[m][:s[s]
a[ XXX]", "yyyy-MM-dd'T']H[H]:m[m][:s[s][.SSS[SSS]]][XXX]", "yyyy-MM-
dd'T'H[H]:m[m][:s[s][.SS
SSSS][XXX]"], "timestamp_with_timezone_format":["yyyy-MM-dd'T'H[H]:m[m][:s[s]
[.SSS[SSS]]][XXX]", "yyyy-MM-dd H[H]:m[m][:s[s][.SSS[SSS]]][XXX]", "yyyy-MM-
dd'T'H[H]
:m[m][:s[s][.SSSSSS][XXX]"], "props":
[{"drop_after_loading":false, "type":"integer", "column":2, "name":"ID", "dimensi
on":0}, {"drop_after_loading":false, "type":
"double", "column":3, "name":"pagerank", "dimension":0}], "format":"pgb", "time_fo
rmat":["h[h]:m[m][:s[s][.SSS] a[ XXX]", "yyyy-MM-dd'T']H[H]:m[m][:s[s]
[.SSS[SSS]]]
[XXX]", "yyyy-MM-dd'T'H[H]:m[m][:s[s][.SSSSSS]
[XXX]"], "key_type":"long", "header":false, "has_keys":true, "error_handling":
{}], "uris":["/scratch/PG/snapshot/snapsho
t_node.pgb"], "label":"node", "key_column":1, "local_date_format":["yyyy-M[M]-
d[d]", "M[M]/d[d]/yyyy", "d[d]-MMM-yyyy", "d[d]-M[M]-yyyy", "yyyy-MM-
dd'T'H[H]:m[m][:s[s]
[.SSS[SSS]]][XXX]", "yyyy-MM-dd'T'H[H]:m[m][:s[s][.SSSSSS]
[XXX]"], "name":"node"}], "timestamp_format":["yyyy-MM-dd'T'H[H]:m[m][:s[s]
[.SSS[SSS]]][XXX]", "yyyy-MM-d
d H[H]:m[m][:s[s][.SSS[SSS]]][XXX]", "yyyy-MM-dd'T'H[H]:m[m][:s[s][.SSSSSS]
[XXX]"], "edge_id_type":"long"}

// Reload from disk
opg4j> var graphFromDisk = session.readGraphWithProperties(storedPgbConfig)
graphFromDisk ==>
PgxGraph[name=bank_graph_partitioned_3,N=1000,E=5001,created=1635790679596]

// Previously computed properties are still part of the graph and can be
queried
opg4j> graphFromDisk.queryPgql("SELECT x.pagerank MATCH(x) LIMIT
5").print().close()
+-----+
| x.pagerank |
+-----+
| 9.748675243734766E-4 |
| 0.004576976478097233 |
| 5.353395438612155E-4 |
| 0.0013043345794813302 |
| 0.001502117014779663 |
+-----+

```

**Note:**

In the case of partitioned graphs, multiple PGB files are being generated, one for each vertex/edge partition in the graph.

4.8 Executing Built-in Algorithms

The in-memory graph server (PGX) contains a set of built-in algorithms that are available as Java APIs.

The following table provides an overview of the available algorithms, grouped by category.



Note:

These algorithms can be invoked through the `Analyst` interface. See the [Analyst Class](#) in Javadoc for more details.

Table 4-5 Overview of Built-In Algorithms

Category	Algorithms
Classic graph algorithms	Prim's Algorithm
Community detection	Conductance Minimization (Soman and Narang Algorithm), Infomap, Label Propagation, Louvain
Connected components	Strongly Connected Components, Weakly Connected Components (WCC)
Link prediction	WTF (Whom To Follow) Algorithm
Matrix factorization	Matrix Factorization
Other	Graph Traversal Algorithms
Path finding	All Vertices and Edges on Filtered Path, Bellman-Ford Algorithms, Bidirectional Dijkstra Algorithms, Compute Distance Index, Compute High-Degree Vertices, Dijkstra Algorithms, Enumerate Simple Paths, Fast Path Finding, Fattest Path, Filtered Fast Path Finding, Hop Distance Algorithms
Ranking and walking	Closeness Centrality Algorithms, Degree Centrality Algorithms, Eigenvector Centrality, Hyperlink-Induced Topic Search (HITS), PageRank Algorithms, Random Walk with Restart, Stochastic Approach for Link-Structure Analysis (SALSA) Algorithms, Vertex Betweenness Centrality Algorithms
Structure evaluation	Adamic-Adar index, Bipartite Check, Conductance, Cycle Detection Algorithms, Degree Distribution Algorithms, Eccentricity Algorithms, K-Core, Local Clustering Coefficient (LCC), Modularity, Partition Conductance, Reachability Algorithms, Topological Ordering Algorithms, Triangle Counting Algorithms

This following topics describe the use of the in-memory graph server (PGX) using Triangle Counting and PageRank analytics as examples.

- [About Built-In Algorithms in the In-Memory Graph Server \(PGX\)](#)
- [Running the Triangle Counting Algorithm](#)
- [Running the PageRank Algorithm](#)

4.8.1 About Built-In Algorithms in the In-Memory Graph Server (PGX)

The in-memory graph server (PGX) contains a set of built-in algorithms that are available as Java APIs. The details of the APIs are documented in the Javadoc that is included in the product documentation library. Specifically, see the `BuiltInAlgorithms` interface Method Summary for a list of the supported in-memory analyst methods.

For example, this is the PageRank procedure signature:

```
/**
 * Classic pagerank algorithm. Time complexity: O(E * K) with E = number
 * of edges, K is a given constant (max
 * iterations)
 *
 * @param graph
 *       graph
 * @param e
 *       maximum error for terminating the iteration
 * @param d
 *       damping factor
 * @param max
 *       maximum number of iterations
 * @return Vertex Property holding the result as a double
 */
public <ID extends Comparable<ID>> VertexProperty<ID, Double>
pagerank(PgxGraph graph, double e, double d, int max);
```

4.8.2 Running the Triangle Counting Algorithm

For triangle counting, the `sortByDegree` boolean parameter of `countTriangles()` allows you to control whether the graph should first be sorted by degree (`true`) or not (`false`). If `true`, more memory will be used, but the algorithm will run faster; however, if your graph is very large, you might want to turn this optimization off to avoid running out of memory.

-
- [JShell](#)
 - [Java](#)

JShell

```
opg4j> analyst.countTriangles(graph, true)
==> 1
```

Java

```
import oracle.pgx.api.*;

Analyst analyst = session.createAnalyst();
long triangles = analyst.countTriangles(graph, true);
```

The algorithm finds one triangle in the sample graph.

 **Tip:**

When using the graph shell, you can increase the amount of log output during execution by changing the logging level. See information about the `:loglevel` command with `:h :loglevel`.

4.8.3 Running the PageRank Algorithm

PageRank computes a rank value between 0 and 1 for each vertex (node) in the graph and stores the values in a `double` property. The algorithm therefore creates a *vertex property* of type `double` for the output.

In the in-memory graph server (PGX), there are two types of vertex and edge properties:

- **Persistent Properties:** Properties that are loaded with the graph from a data source are fixed, in-memory copies of the data on disk, and are therefore persistent. Persistent properties are read-only, immutable and shared between sessions.
- **Transient Properties:** Values can only be written to transient properties, which are private to a session. You can create transient properties by calling `createVertexProperty` and `createEdgeProperty` on `PgxGraph` objects, or by copying existing properties using `clone()` on `Property` objects.

Transient properties hold the results of computation by algorithms. For example, the PageRank algorithm computes a rank value between 0 and 1 for each vertex in the graph and stores these values in a transient property named `pg_rank`. Transient properties are destroyed when the `Analyst` object is destroyed.

This example obtains the top three vertices with the highest PageRank values. It uses a transient vertex property of type `double` to hold the computed PageRank values. The PageRank algorithm uses the following default values for the input parameters: error (tolerance = 0.001), damping factor = 0.85, and maximum number of iterations = 100.

-
- [JShell](#)
 - [Java](#)

JShell

```
opg4j> rank = analyst.pagerank(graph, 0.001, 0.85, 100);  
==> ...  
opg4j> rank.getTopKValues(3)  
==> 128=0.1402019732468347  
==> 333=0.12002296283541904  
==> 99=0.09708583862990475
```

Java

```
import java.util.Map.Entry;
import oracle.pgx.api.*;

Analyst analyst = session.createAnalyst();
VertexProperty<Integer, Double> rank = analyst.pagerank(graph, 0.001, 0.85,
100);
for (Entry<Integer, Double> entry : rank.getTopKValues(3)) {
    System.out.println(entry.getKey() + "=" + entry.getValue());
}
```

4.9 Using Custom PGX Graph Algorithms

A custom PGX graph algorithm allows you to write a graph algorithm in Java syntax and have it automatically compiled to an efficient parallel implementation.

For more detailed information that appears in the following subtopics, see the [PGX Algorithm Specification](#).

- [Writing a Custom PGX Algorithm](#)
- [Compiling and Running a Custom PGX Algorithm](#)
- [Example Custom PGX Algorithm: PageRank](#)

4.9.1 Writing a Custom PGX Algorithm

A PGX algorithm is a regular .java file with a single class definition that is annotated with `@GraphAlgorithm`. For example:

```
import oracle.pgx.algorithm.annotations.GraphAlgorithm;

@GraphAlgorithm
public class MyAlgorithm {
    ...
}
```

A PGX algorithm class must contain exactly one public method which will be used as entry point. The class may contain any number of private methods.

For example:

```
import oracle.pgx.algorithm.PgxGraph;
import oracle.pgx.algorithm.VertexProperty;
import oracle.pgx.algorithm.annotations.GraphAlgorithm;
import oracle.pgx.algorithm.annotations.Out;

@GraphAlgorithm
public class MyAlgorithm {
    public int myAlgorithm(PgxGraph g, @Out VertexProperty<Integer>
distance) {
```

```
        System.out.println("My first PGX Algorithm program!");  
  
        return 42;  
    }  
}
```

As with normal Java methods, a PGX algorithm method only supports primitive data types as return values (an integer in this example). More interesting is the `@Out` annotation, which marks the vertex property `distance` as output parameter. The caller passes output parameters by reference. This way, the caller has a reference to the modified property after the algorithm terminates.

- [Collections](#)
- [Iteration](#)
- [Reductions](#)

4.9.1.1 Collections

To create a collection you call the `.create()` function. For example, a `VertexProperty<Integer>` is created as follows:

```
VertexProperty<Integer> distance = VertexProperty.create();
```

To get the value of a property at a certain vertex `v`:

```
distance.get(v);
```

Similarly, to set the property of a certain vertex `v` to a value `e`:

```
distance.set(v, e);
```

You can even create properties of collections:

```
VertexProperty<VertexSequence> path = VertexProperty.create();
```

However, PGX Algorithm assignments are always *by value* (as opposed to *by reference*). To make this explicit, you *must* call `.clone()` when assigning a collection:

```
VertexSequence sequence = path.get(v).clone();
```

Another consequence of values being passed *by value* is that you can check for equality using the `==` operator instead of the Java method `.equals()`. For example:

```
PgxVertex v1 = G.getRandomVertex();  
PgxVertex v2 = G.getRandomVertex();  
System.out.println(v1 == v2);
```

4.9.1.2 Iteration

The most common operations in PGX algorithms are iterations (such as looping over all vertices, and looping over a vertex's neighbors) and graph traversal (such as breath-first/depth-first). All collections expose a `forEach` and `forSequential` method by which you can iterate over the collection in parallel and in sequence, respectively.

For example:

- To iterate over a graph's vertices in parallel:

```
G.getVertices().forEach(v -> {  
    ...  
});
```

- To iterate over a graph's vertices in sequence:

```
G.getVertices().forSequential(v -> {  
    ...  
});
```

- To traverse a graph's vertices from `r` in breadth-first order:

```
import oracle.pgx.algorithm.Traversal;  
  
Traversal.inBFS(G, r).forward(n -> {  
    ...  
});
```

Inside the `forward` (or `backward`) lambda you can access the current level of the BFS (or DFS) traversal by calling `currentLevel()`.

4.9.1.3 Reductions

Within these parallel blocks it is common to atomically update, or reduce to, a variable defined outside the lambda. These atomic reductions are available as methods on `Scalar<T>`: `reduceAdd`, `reduceMul`, `reduceAnd`, and so on. For example, to count the number of vertices in a graph:

```
public int countVertices() {  
    Scalar<Integer> count = Scalar.create(0);  
  
    G.getVertices().forEach(n -> {  
        count.reduceAdd(1);  
    });  
  
    return count.get();  
}
```

Sometimes you want to update multiple values atomically. For example, you might want to find the smallest property value as well as the vertex whose property value attains this

smallest value. Due to the parallel execution, two separate reduction statements might get you in an inconsistent state.

To solve this problem the `Reductions` class provides `argMin` and `argMax` functions. The first argument to `argMin` is the current value and the second argument is the potential new minimum. Additionally, you can chain `andUpdate` calls on the `ArgMinMax` object to indicate other variables and the values that they should be updated to (atomically). For example:

```
VertexProperty<Integer> rank = VertexProperty.create();
int minRank = Integer.MAX_VALUE;
PgxVertex minVertex = PgxVertex.NONE;

G.getVertices().forEach(n ->
    argMin(minRank, rank.get(n)).andUpdate(minVertex, n)
);
```

4.9.2 Compiling and Running a Custom PGX Algorithm

To be able to compile and run a custom PGX algorithm, you must perform the following actions:

1. Set the following two configuration parameters in the `conf/pgx.conf` file:
 - Set the `graph_algorithm_language` option to `JAVA`.
 - Set the `java_home_dir` option to the path to your Java home (use `<system-java-home-dir>` to have PGX infer Java home from the system properties).

```
{
  "graph_algorithm_language": "JAVA",
  "java_home_dir": "<system-java-home-dir>"
}
```

2. Create a session.

-
- [JShell](#)
 - [Java](#)
 - [Python](#)

JShell

```
cd /opt/oracle/graph
./bin/opg4j
```

Java

```
import oracle.pgx.algorithm.*;
PgxSession session = Pgx.createSession("my-session");
```

Python

```
session = instance.create_session("my-session")
```

3. Compile a PGX Algorithm. For example:
-

- [JShell](#)
- [Java](#)
- [Python](#)

JShell

```
opg4j> var myAlgorithm = session.compileProgram("/path/to/  
MyAlgorithm.java")  
myAlgorithm ==> CompiledProgram[name=MyAlgorithm]
```

Java

```
import oracle.pgx.algorithm.CompiledProgram;  
CompiledProgram myAlgorithm = session.compileProgram("/path/to/  
MyAlgorithm.java");
```

Python

```
my_algorithm = session.compile_program("/path/to/MyAlgorithm.java")
```

4. Run the algorithm. For example:
-

- [JShell](#)
- [Java](#)
- [Python](#)

JShell

```
opg4j> var graph =session.readGraphWithProperties("/path/to/  
bank_graph_analytics.json")  
graph ==>  
PgxGraph[name=bank_graph_analytics,N=1000,E=5001,created=1633504705054]  
opg4j> var property = graph.createVertexProperty(PropertyType.INTEGER)  
property ==>
```

```
VertexProperty[name=vertex_prop_integer_9,type=integer,graph=bank_graph_analytics]
opg4j> myAlgorithm.run(graph, property)
$6 ==> {
  "success" : true,
  "canceled" : false,
  "exception" : null,
  "returnValue" : 42,
  "executionTimeMs" : 0
}
```

Java

```
import oracle.pgx.algorithm.VertexProperty;
PgxGraph graph = session.readGraphWithProperties("/path/to/bank_graph_analytics.json");
VertexProperty property =
graph.createVertexProperty(PropertyType.INTEGER);
myAlgorithm.run(graph, property);
```

Python

```
graph = session.read_graph_with_properties("/path/to/bank_graph_analytics.json")
property = graph.create_vertex_property("integer")
my_algorithm.run(graph, property)
{'success': True, 'canceled': False, 'exception': None, 'return_value': 42, 'execution_time(ms)': 1}
```

4.9.3 Example Custom PGX Algorithm: PageRank

The following is an implementation of pagerank as a PGX algorithm:

```
import oracle.pgx.algorithm.PgxGraph;
import oracle.pgx.algorithm.Scalar;
import oracle.pgx.algorithm.VertexProperty;
import oracle.pgx.algorithm.annotations.GraphAlgorithm;
import oracle.pgx.algorithm.annotations.Out;

@GraphAlgorithm
public class Pagerank {
    public void pagerank(PgxGraph G, double tol, double damp, int
max_iter, boolean norm, @Out VertexProperty<Double> rank) {
        Scalar<Double> diff = Scalar.create();
        int cnt = 0;
        double N = G.getNumVertices();

        rank.setAll(1 / N);
        do {
            diff.set(0.0);
```



```
Scalar<Double> dangling_factor = Scalar.create(0d);

if (norm) {
    dangling_factor.set(damp / N * G.getVertices().filter(v ->
v.getOutDegree() == 0).sum(rank::get));
}

G.getVertices().forEach(t -> {
    double in_sum = t.getInNeighbors().sum(w -> rank.get(w) /
w.getOutDegree());
    double val = (1 - damp) / N + damp * in_sum + dangling_factor.get();
    diff.reduceAdd(Math.abs(val - rank.get(t)));
    rank.setDeferred(t, val);
});
cnt++;
} while (diff.get() > tol && cnt < max_iter);
}
```

4.10 Creating Subgraphs

You can create subgraphs based on a graph that has been loaded into memory. You can use filter expressions or create bipartite subgraphs based on a vertex (node) collection that specifies the left set of the bipartite graph.

For information about reading a graph into memory, see [Loading a Graph Into the Graph Server \(PGX\)](#) for the various methods to load a graph into the in-memory graph server (PGX).


- [About Filter Expressions](#)
- [Using a Simple Filter to Create a Subgraph](#)
- [Using a Complex Filter to Create a Subgraph](#)
- [Using a Vertex Set to Create a Bipartite Subgraph](#)

4.10.1 About Filter Expressions

Filter expressions are expressions that are evaluated for each vertex or edge. The expression can define predicates that a vertex or an edge must fulfil in order to be contained in the result, in this case a subgraph.

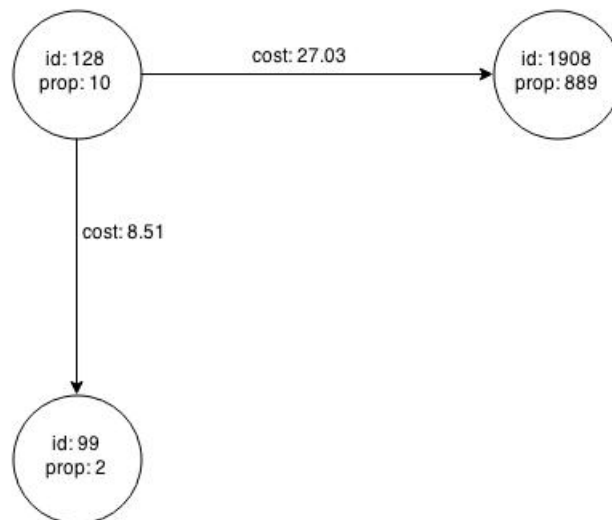
Consider an example graph that consists of four vertices (nodes) and four edges. For an edge to match the filter expression `src.prop == 10`, the source vertex `prop` property must equal 10. Two edges match that filter expression, as shown in the following figure.

Figure 4-5 Edges Matching `src.prop == 10`



The following figure shows the graph that results when the filter is applied.

Figure 4-6 Graph Created by the Simple Filter



The vertex filter `src.prop == 10` filters out the edges associated with vertex 333 and the vertex itself.

4.10.2 Using a Simple Filter to Create a Subgraph

The following examples create the subgraph described in [About Filter Expressions](#).

- [JShell](#)
- [Java](#)

JShell

```
var subgraph = graph.filter(new VertexFilter("vertex.prop == 10"))
```

Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.filter.*;

PgxGraph graph = session.readGraphWithProperties(...);
PgxGraph subgraph = graph.filter(new VertexFilter("vertex.prop == 10"));
```


4.10.3 Using a Complex Filter to Create a Subgraph

This example uses a slightly more complex filter. It uses the `outDegree` function, which calculates the number of outgoing edges for an identifier (source `src` or destination `dst`). The following filter expression matches all edges with a `cost` property value greater than 50 and a destination vertex (node) with an `outDegree` greater than 1.


```
dst.outDegree() > 1 && edge.cost > 50
```

One edge in the sample graph matches this filter expression, as shown in the following figure.

Figure 4-7 Edges Matching the `outDegree` Filter




The following figure shows the graph that results when the filter is applied. The filter excludes the edges associated with the vertices 99 and 1908, and so excludes those vertices also.

Figure 4-8 Graph Created by the outDegree Filter

4.10.4 Using a Vertex Set to Create a Bipartite Subgraph

You can create a bipartite subgraph by specifying a set of vertices (nodes), which are used as the left side. A bipartite subgraph has edges only between the left set of vertices and the right set of vertices. There are no edges within those sets, such as between two nodes on the left side. In the in-memory graph server (PGX), vertices that are isolated because all incoming and outgoing edges were deleted are not part of the bipartite subgraph.

The following figure shows a bipartite subgraph. No properties are shown.



The following examples create a bipartite subgraph from a simple graph consisting of four vertices and four edges. The vertex ID values for the four vertices are 99, 128, 1908 and 333 respectively. See [Figure 4-5 in About Filter Expressions](#) for more information on the vertex and edge property values including the edge direction between the vertices.

You must first create a vertex collection and fill it with the vertices for the left side. In the example shown, vertices with vertex ID values 333 and 99 are added to the left side of the vertex collection.

Using the Shell to Create a Bipartite Subgraph

```
opg4j> s = graph.createVertexSet()
==> ...
opg4j> s.addAll([graph.getVertex(333), graph.getVertex(99)])
==> ...
opg4j> s.size()
==> 2
opg4j> bGraph = graph.bipartiteSubGraphFromLeftSet(s)
==> PGX Bipartite Graph named sample-sub-graph-4
```


Using Java to Create a Bipartite Subgraph

```
import oracle.pgx.api.*;

VertexSet<Integer> s = graph.createVertexSet();
s.addAll(graph.getVertex(333), graph.getVertex(99));
BipartiteGraph bGraph = graph.bipartiteSubGraphFromLeftSet(s);
```

When you create a subgraph, the in-memory graph server (PGX) automatically creates a Boolean vertex (node) property that indicates whether the vertex is on the left side. You can specify a unique name for the property.

The resulting bipartite subgraph looks like this:



Vertex with ID 1908 is excluded from the bipartite subgraph. The only edge that connected that vertex extended from 128 to 1908. The edge was removed, because it violated the bipartite properties of the subgraph. Vertex 1908 had no other edges, and so was removed as well. Moreover, the edge from the vertex with the ID 128 to the vertex with ID 99 is not present in the bipartite subgraph, because edges are only allowed to go from left to right (and not from right to left).

4.11 Using Automatic Delta Refresh to Handle Database Changes

You can automatically refresh (auto-refresh) graphs periodically to keep the in-memory graph synchronized with changes to the property graph stored in the property graph tables in Oracle Database (VT\$ and GE\$ tables).

Note that the auto-refresh feature is not supported when loading a graph into PGX in memory directly from relational tables.

- [Configuring the In-Memory Server for Auto-Refresh](#)
- [Configuring Basic Auto-Refresh](#)
- [Reading the Graph Using the In-Memory Graph Server \(PGX\) or a Java Application](#)
- [Checking Out a Specific Snapshot of the Graph](#)
- [Advanced Auto-Refresh Configuration](#)
- [Special Considerations When Using Auto-Refresh](#)

4.11.1 Configuring the In-Memory Server for Auto-Refresh

Because auto-refresh can create many snapshots and therefore may lead to a high memory usage, by default the option to enable auto-refresh for graphs is available only to administrators.

To allow all users to auto-refresh graphs, you must include the following line into the in-memory graph server (PGX) configuration file (located in `$ORACLE_HOME/md/property_graph/pgx/conf/pgx.conf`):

```
{
  "allow_user_auto_refresh": true
}
```

4.11.2 Configuring Basic Auto-Refresh

Auto-refresh is configured in the loading section of the graph configuration. The example in this topic sets up auto-refresh to check for updates every minute, and to create a new snapshot when the data source has changed.

The following block (JSON format) enables the auto-refresh feature in the configuration file of the sample graph:

```
{
  "format": "pg",
  "jdbc_url": "jdbc:oracle:thin:@mydatabaseserver:1521/dbName",
  "username": "scott",
  "password": "<password>",
  "name": "my_graph",
  "vertex_props": [{
    "name": "prop",
    "type": "integer"
  }],
  "edge_props": [{
    "name": "cost",
    "type": "double"
  }],
  "separator": " ",
  "loading": {
    "auto_refresh": true,
    "update_interval_sec": 60
  },
}
```

Notice the additional `loading` section containing the auto-refresh settings. You can also use the Java APIs to construct the same graph configuration programmatically:

```
GraphConfig config = GraphConfigBuilder.forPropertyGraphRdbms()
    .setJdbcUrl("jdbc:oracle:thin:@mydatabaseserver:1521/dbName")
    .setUsername("scott")
    .setPassword("<password>")
    .setName("my_graph")
    .addVertexProperty("prop", PropertyType.INTEGER)
    .addEdgeProperty("cost", PropertyType.DOUBLE)
    .setAutoRefresh(true)
    .setUpdateIntervalSec(60)
    .build();
```

4.11.3 Reading the Graph Using the In-Memory Graph Server (PGX) or a Java Application

After creating the graph configuration, you can load the graph into the in-memory graph server (PGX) using the regular APIs.

```
opg4j> G = session.readGraphWithProperties("graphs/my-config.pg.json")
```

After the graph is loaded, a background task is started automatically, and it periodically checks the data source for updates.

4.11.4 Checking Out a Specific Snapshot of the Graph

The database is queried every minute for updates. If the graph has changed in the database after the time interval passed, the graph is reloaded and a new snapshot is created in-memory automatically.

You can "check out" (move a pointer to a different version of) the available in-memory snapshots of the graph using the `getAvailableSnapshots()` method of `PgxSession`. Example output is as follows:

```
opg4j> session.getAvailableSnapshots(G)
==> GraphMetaData [getNumVertices()=4, getNumEdges()=4, memoryMb=0,
dataSourceVersion=1453315103000,
creationRequestTimestamp=1453315122669 (2016-01-20 10:38:42.669),
creationTimestamp=1453315122685 (2016-01-20 10:38:42.685),
vertexIdType=integer, edgeIdType=long]
==> GraphMetaData [getNumVertices()=5, getNumEdges()=5, memoryMb=3,
dataSourceVersion=1452083654000,
creationRequestTimestamp=1453314938744 (2016-01-20 10:35:38.744),
creationTimestamp=1453314938833 (2016-01-20 10:35:38.833),
vertexIdType=integer, edgeIdType=long]
```

The preceding example output contains two entries, one for the originally loaded graph with 4 vertices and 4 edges, and one for the graph created by auto-refresh with 5 vertices and 5 edges.

To check out a specific snapshot of the graph, use the `setSnapshot()` methods of `PgxSession` and give it the `creationTimestamp` of the snapshot you want to load.

For example, if `G` is pointing to the newer graph with 5 vertices and 5 edges, but you want to analyze the older version of the graph, you need to set the snapshot to 1453315122685. In the graph shell:

```
opg4j> G.getNumVertices()
==> 5
opg4j> G.getNumEdges()
==> 5

opg4j> session.setSnapshot(G, 1453315122685)
==> null
```



```
opg4j> G.getNumVertices()
==> 4
opg4j> G.getNumEdges()
==> 4
```

You can also load a specific snapshot of a graph directly using the `readGraphAsOf()` method of `PgxSession`. This is a shortcut for loading a graph with `readGraphWithProperty()` followed by a `setSnapshot()`. For example:

```
opg4j> G = session.readGraphAsOf( config, 1453315122685 )
```

If you do not know or care about what snapshots are currently available in-memory, you can also specify a time span of how “old” a snapshot is acceptable by specifying a maximum allowed age. For example, to specify a maximum snapshot age of 60 minutes, you can use the following:

```
opg4j> G = session.readGraphWithProperties( config, 60l, TimeUnit.MINUTES )
```

If there are one or more snapshots in memory younger (newer) than the specified maximum age, the youngest (newest) of those snapshots will be returned. If all the available snapshots are older than the specified maximum age, or if there is no snapshot available at all, then a new snapshot will be created automatically.

4.11.5 Advanced Auto-Refresh Configuration

You can specify advanced options for auto-refresh configuration.

Internally, the in-memory graph server (PGX) fetches the changes since the last check from the database and creates a new snapshot by applying the delta (changes) to the previous snapshot. There are two timers: one for fetching and caching the deltas from the database, the other for actually applying the deltas and creating a new snapshot.

Additionally, you can specify a threshold for the number of cached deltas. If the number of cached changes grows above this threshold, a new snapshot is created automatically. The number of cached changes is a simple sum of the number of vertex changes plus the number of edge changes.

The deltas are fetched periodically and cached on the in-memory graph server for two reasons:

- To speed up the actual snapshot creation process
- To account for the case that the database can “forget” changes after a while

You can specify both a threshold and an update timer, which means that both conditions will be checked before new snapshot is created. At least one of these parameters (threshold or update timer) must be specified to prevent the delta cache from becoming too large. The interval at which the source is queried for changes must not be omitted.

The following parameters show a configuration where the data source is queried for new deltas every 5 minutes. New snapshots are created every 20 minutes or if the cached deltas reach a size of 1000 changes.

```
{
  "format": "pg",
```

```
"jdbc_url": "jdbc:oracle:thin:@mydatabaseserver:1521/dbName",
"username": "scott",
"password": "<your_password>",
"name": "my_graph",

"loading": {
  "auto_refresh": true,
  "fetch_interval_sec": 300,
  "update_interval_sec": 1200,
  "update_threshold": 1000,
  "create_edge_id_index": true,
  "create_edge_id_mapping": true
}
}
```

4.11.6 Special Considerations When Using Auto-Refresh

This section explains a few special considerations when you enable auto-refresh for graphs in the in-memory graph server (PGX):

- If you call `graph.destroy()`, auto-refresh does not immediately stop. It only stops once the graph is actually freed from the server memory. This happens when all the following conditions are true:
 1. No other session is referencing that graph.
 2. PGX consumes more than `release_memory_threshold` memory. `release_memory_threshold` is a `pgx.conf` option that defaults to 85% of available system memory.
 3. The PGX "garbage collector" has been run. `memory_cleanup_interval` is a `pgx.conf` option which defaults to once every 10 minutes.
- If you configure the graph to be loaded with auto-refresh, you cannot omit the `jdbc_url`, `username` and `keystore` parameters from the graph configuration file since auto-refreshed graphs are not "user bound". You cannot obtain the connection settings from the user who initiated it.

4.12 Starting the In-Memory Graph Server (PGX)

This section describes the commands to start and stop the in-memory graph server (PGX).

A preconfigured version of Apache Tomcat is bundled, which allows you to start the in-memory graph server (PGX) by running a script.

As a prerequisite to start the graph server in remote mode, you must ensure that Oracle graph server is installed in your system. See [Installing Oracle Graph Server](#) for instructions to install the graph server (PGX).

 **Note:**

See [Usage Modes of the In-memory Graph Server \(PGX\)](#) for more information on the different graph server execution modes.

- [Starting and Stopping the Graph Server \(PGX\) Using the Command Line](#)
- [Configuring the In-Memory Graph Server \(PGX\)](#)

4.12.1 Starting and Stopping the Graph Server (PGX) Using the Command Line

PGX is integrated with `systemd` to run it as a Linux service in the background.

If you need to configure the server before starting it, see [Configuring the In-Memory Graph Server \(PGX\)](#) and [Configuration Parameters for the Graph Server \(PGX\) Engine](#) for more information on the configuration options.

The commands to start and stop the graph server (PGX) and the PGX engine are as follows:

 **Note:**

You can run the following commands without `sudo` if you are the root user.

To start the PGX server as a daemon process, run the following command:

```
sudo systemctl start pgx
```

To stop the server, run the following command:

```
sudo systemctl stop pgx
```

If the server does not start up, you can see if there are any errors by running:

```
sudo journalctl -u pgx.service
```

For more information about how to interact with `systemd` on Oracle Linux, see the Oracle Linux administrator's documentation.

4.12.2 Configuring the In-Memory Graph Server (PGX)

You can configure the in-memory graph server (PGX) by modifying the `/etc/oracle/graph/server.conf` file. The following table shows the valid configuration options, which can be specified in JSON format.

Table 4-6 Configuration Parameters for the In-Memory Graph Server (PGX)

Parameter	Type	Description	Default
ca_certs	array of string	List of files storing trusted certificates (PEM format). If <code>enable_tls</code> is set to <code>false</code> , this field has no effect.	[]

Table 4-6 (Cont.) Configuration Parameters for the In-Memory Graph Server (PGX)

Parameter	Type	Description	Default
ciphers	array of string	List of cipher suites to be used by the server. For example, [cipher1, cipher2.]	["TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256", "TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384", "TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256", "TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384", "TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256", "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256", "TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384", "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384", "TLS_DHE_RSA_WITH_AES_128_GCM_SHA256", "TLS_DHE_DSS_WITH_AES_128_GCM_SHA256", "TLS_DHE_RSA_WITH_AES_128_CBC_SHA256", "TLS_DHE_DSS_WITH_AES_128_CBC_SHA256", "TLS_DHE_DSS_WITH_AES_256_GCM_SHA384", "TLS_DHE_RSA_WITH_AES_256_CBC_SHA256", "TLS_DHE_DSS_WITH_AES_256_CBC_SHA256", "TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA", "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA", "TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA", "TLS_DHE_DSS_WITH_AES_128_CBC_SHA", "TLS_DHE_RSA_WITH_AES_128_CBC_SHA", "TLS_DHE_DSS_WITH_AES_256_CBC_SHA", "TLS_DHE_RSA_WITH_AES_256_CBC_SHA", "TLS_RSA_WITH_AES_128_GCM_SHA256", "TLS_DH_DSS_WITH_AES_128_GCM_SHA256", "TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256",

Table 4-6 (Cont.) Configuration Parameters for the In-Memory Graph Server (PGX)

Parameter	Type	Description	Default
			<pre> "TLS_RSA_WITH_AES_256_GCM_SHA384", "TLS_DH_DSS_WITH_AES_256_GCM_SHA384", "TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384", "TLS_RSA_WITH_AES_128_CBC_SHA256", "TLS_DH_DSS_WITH_AES_128_CBC_SHA256", "TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256", "TLS_RSA_WITH_AES_256_CBC_SHA256", "TLS_DH_DSS_WITH_AES_256_CBC_SHA256", "TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384", "TLS_RSA_WITH_AES_128_CBC_SHA", "TLS_DH_DSS_WITH_AES_128_CBC_SHA", "TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA", "TLS_RSA_WITH_AES_256_CBC_SHA", "TLS_DH_DSS_WITH_AES_256_CBC_SHA", "TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA"] </pre>
context_path	string	This can be used to change the context path. For example, if you specify port as 7007 and context path as /pgx, the server will listen on https://localhost:7007/pgx	/
enable_tls	boolean	If true, the server enables transport layer security (TLS).	true
port	integer	Port the graph server (PGX) server should listen on.	7007

Table 4-6 (Cont.) Configuration Parameters for the In-Memory Graph Server (PGX)

Parameter	Type	Description	Default
server_certificate	string	The path to the server certificate to be presented to TLS clients (PEM format).	NULL


 **Note:** This file must only contain one certificate

Table 4-6 (Cont.) Configuration Parameters for the In-Memory Graph Server (PGX)

Parameter	Type	Description	Default
			. l f y o u r c e r t i f i c a t e i s a c h a i n a n d c o n t a i n s a r o o t c e r t i f i c a t

Table 4-6 (Cont.) Configuration Parameters for the In-Memory Graph Server (PGX)

Parameter	Type	Description	Default
			enable_certificate_signing
server_private_key	string	If <code>enable_tls</code> is set to <code>false</code> , this field has no effect This is the path to the file storing the private key of the server (PEM format). For security reasons, the file must have only Read and Write permissions only for the owner (600 permissions in a POSIX filesystem), otherwise an error will be thrown. If <code>enable_tls</code> is set to <code>false</code> , this field has no effect	NULL
tls_version	string	TLS version to be used by the server. For example, TLSv1.2	TLSv1.2

Table 4-6 (Cont.) Configuration Parameters for the In-Memory Graph Server (PGX)

Parameter	Type	Description	Default
<code>working_directory</code>	string	The working directory used by the server to store temporary files. Needs to be writeable by the process which started the server and should not be touched by any other process while the server is running.	

The in-memory graph server (PGX) enables two-way SSL/TLS (Transport Layer Security) by default. The server enforces TLS 1.2 and disables certain cipher suites known to be vulnerable to attacks. Upon a TLS handshake, both the server and the client present certificates to each other, which are used to validate the authenticity of the other party. Client certificates are also used to authorize client applications.

The following is an example `server.conf` configuration file:

```
{
  "port": 7007,
  "enable_tls": true,
  "server_cert": "server_cert.pem",
  "server_private_key": "server_key.pem",
  "ca_certs": [
    "server_cert.pem"
  ]
}
```

4.13 Connecting to the In-Memory Graph Server (PGX)

This section explains how to connect to the in-memory graph server (PGX) running in remote mode or when deployed as a web application on Apache Tomcat or Oracle WebLogic Server.

The prerequisite requirement to connect to the graph server is to have the in-memory graph server (PGX) up and running. See [Starting and Stopping the Graph Server \(PGX\) Using the Command Line](#) for more information on the commands to start the graph server.

Note:

If you are using the graph server (PGX) as a library, see [Using Graph Server \(PGX\) as a Library](#) for more information.

- [Connecting with the Graph Shell](#)
- [Connecting with Java](#)
- [Connecting with Python](#)

4.13.1 Connecting with the Graph Shell

The simplest way to connect to a remote graph server (PGX) instance is to specify the base URL of the server along with the database user name required for the graph server (PGX) authentication as shown:

- [JShell](#)
- [Python](#)

JShell

```
cd /opt/oracle/graph
./bin/opg4j --base_url https://<host>:<port> --username <graphuser>
```

Python

```
cd /opt/oracle/graph
./bin/opg4py --base_url https://<host>:<port> --username <graphuser>
```

where :

- `<host>`: is the server host name
- `<port>`: is the server port
- `<graphuser>`: is the database user

 **Note:**

You will be prompted for the database password.

See [User Authentication and Authorization](#) for more information.

About Logging HTTP Requests

The graph shell suppresses all debugging messages by default. To see which HTTP requests are executed, set the log level for `oracle.pgx` to `DEBUG`, as shown in this example:

 **Note:**

Enabling these logs can lead to sensitive information like passwords getting printed on the screen.

- JShell
- Python

JShell

```

opg4j> loglevel("oracle.pgx","DEBUG")
==> Log level of oracle.pgx logger set to DEBUG
opg4j> session.readGraphWithProperties("bank_graph_analytics.json",
"bank_graph");
06:29:03,702 DEBUG CommonsVfsProvider - resolve
bank_graph_analytics.json
06:29:03,702 DEBUG AbstractConfigFactory - parse graph config from
bank_graph_analytics.json (parent: file:///opt/oracle/graph)
06:29:03,709 DEBUG RemoteUtils - create session cookie (session ID =
f5d029d7-2924-4cd4-86a9-6999c1ce5e3f)
06:29:03,713 DEBUG RemoteUtils - no value for the sticky cookie given
06:29:03,713 DEBUG RemoteUtils - create csrf token cookie (token =
36acbee2-6b78-4c13-b114-41040809833a)
06:29:03,713 DEBUG HttpRequestExecutor - Requesting POST https://
localhost:7007/core/v1/loadGraph HTTP/1.1 with payload
{"graphConfig":"HRcBVFVcO0dfXU9bUEhGEEYodkMUElZYRFpcZgBeDxBYcxFcRGY2c21
wIBUBEElAW11HQV5vVhpYAFRIFAMbeC5+NithRx9EEkwUVRADRLFBXVhGFlBpT0ZfSEFpAl
ZBQ1RXG1kTKiEXSREVCUAWU1dmElJfRlXKallGDBJdSV1EQwMpd1paVhZfFXYXSREIBlgBE
ggbMEVMXEpUUV9HQUGwSV1FFVBDV01QVgluAApZEF4IRA9GdHdqMGhkdhseFklREBBdQ1lC
CFZDaU9cSxdUGzpFF1wQD1EBQhADRnZOUVZHW11HQUGwQVdXBVBDURsDQkFSEQBUEVY5DVA
db19YFEdEXF4QDktVDxdRUBQUBVhZV1tYSgZuFwRXCVY5CFQJVRADRNvsfHJtcWlZJjdrbH
ViQxUPXVxAzhdIEwAXXxEKCVsDeh4bAlhfX1hGFhCWEQBWQEsUHGQBFE9cSxdUGzpFF1wQD
1EBQkEbXmxWEFJXTXJXDahBQFYUWxtkchsVGw1QDgAXXxEnBVYLRVxNFxUBEFVdVUldQMW
F0MUAKtIV01cZghUGjpyBEMWD1sDEghNFkJITxUQUEXAAgZVX11pFVhPwLxmVwJcBkcPR3E
nKH47fn19IWQPhtZUVRrFx1ESBoMQ1BDQlxeXBETT0dTckELB0FGChBLAFVAQRtPaQEWDQ
VZSBoMQ1tMWFJmXhFQEWlqBF0HCkwQWVFKRkotMjkyAElr0w==","graphName":"bank_g
raph","_csrf_token":"36acbee2-6b78-4c13-b114-41040809833a"}
06:29:03,788 DEBUG HttpRequestExecutor - received HTTP status 202
06:29:03,789 DEBUG HttpRequestExecutor -
{"futureId":"7f7a2206-8881-4c1e-909f-6e8778be617c"}
06:29:03,789 DEBUG PgxRemoteFuture - Requesting GET https://
localhost:7007/core/v1/futures/x-future-id/status HTTP/1.1
06:29:03,801 DEBUG PgxRemoteFuture - Requesting GET https://
localhost:7007/core/v1/futures/x-future-id/value HTTP/1.1
06:29:03,831 DEBUG RemoteUtils - received HTTP status 201
06:29:03,831 DEBUG RemoteUtils - {"id":"8B473228-0751-49A9-
A945-9A0E4011AB69","links":[{"href":"https://localhost:7007/core/v1/
graphs/x-graph-id","rel":"self","method":"GET","interaction":["async-
polling"]}, {"href":"https://localhost:7007/core/v1/graphs/x-graph-
id","rel":"canonical","method":"GET","interaction":["async-
polling"]}],"graphName":"bank_graph","vertexTables":{"Accounts":
{"name":"Accounts","metaData":
{"name":"Accounts","idType":"integer","labels":
["Accounts"],"properties":[],"edgeProviderNamesWhereSource":
["Transfers"],"edgeProviderNamesWhereDestination":
["Transfers"],"id":null,"links":null},"providerLabels":
["Accounts"],"entityKeyType":"integer","isIdentityKeyMapping":false,"ve
rtexProperties":{},"vertexLabels":{"id":"04156FFE-
A3C1-4A6D-87E5-879A0895BBD4","links":[{"href":"https://localhost:7007/

```

```

core/v1/graphs/x-graph-id/properties/x-property-
name","rel":"self","method":"GET","interaction":["async-polling"]},
{"href":"https://localhost:7007/core/v1/graphs/x-graph-id/properties/x-
property-name","rel":"canonical","method":"GET","interaction":["async-
polling"]}],{"dimension":-1,"propertyId":"04156FFE-
A3C1-4A6D-87E5-879A0895BBD4","name":"__vertex_labels__","entityType":"vertex"
,"type":"ro_string_set","namespace":"2C17C639-3771-3E30-88AE-34D6B380C5EC","t
ransient":false},"transient":false}},{"edgeTables":{"Transfers":
{"name":"Transfers","metaData":
{"name":"Transfers","idType":"long","directed":true,"labels":
["Transfers"],"properties":
[{"name":"AMOUNT","id":null,"propertyType":"float","dimension":0,"transient":
true,"links":null,"propertyId":"AF2A2D0A-9C8C-478F-
BD74-3444A7DD7339"}],"sourceVertexProviderName":"Accounts","destinationVertex
ProviderName":"Accounts","id":null,"links":null},"providerLabels":
["Transfers"],"entityKeyType":"long","isIdentityKeyMapping":true,"sourceVerte
xTableName":"Accounts","destinationVertexTableName":"Accounts","edgePropertie
s":{"4046D845-D0C6-4231-A69B-F69D4963CD91":{"id":"4046D845-D0C6-4231-A69B-
F69D4963CD91","links":[{"href":"https://localhost:7007/core/v1/graphs/x-
graph-id/properties/x-property-
name","rel":"self","method":"GET","interaction":["async-polling"]},
{"href":"https://localhost:7007/core/v1/graphs/x-graph-id/properties/x-
property-name","rel":"canonical","method":"GET","interaction":["async-
polling"]}],{"dimension":0,"propertyId":"4046D845-D0C6-4231-A69B-
F69D4963CD91","name":"AMOUNT","entityType":"edge","type":"float","namespace":
"2C17C639-3771-3E30-88AE-34D6B380C5EC","transient":false},"edgeLabel":
{"id":"9763546A-1860-49A4-9292-77D2AA04F4BB","links
06:29:03,836 DEBUG PgxSession - engine reports latest snapshot is 621849
milli-seconds old. Max age is 0 milli-seconds
06:29:03,836 DEBUG PgxSession - ==> try to check out newer snapshot
06:29:03,836 DEBUG RemoteUtils - create session cookie (session ID =
f5d029d7-2924-4cd4-86a9-6999c1ce5e3f)
06:29:03,836 DEBUG RemoteUtils - no value for the sticky cookie given
06:29:03,836 DEBUG RemoteUtils - create csrf token cookie (token =
36acbee2-6b78-4c13-b114-41040809833a)
06:29:03,836 DEBUG HttpRequestExecutor - Requesting POST https://
localhost:7007/core/v1/graphs/x-graph-id/refresh HTTP/1.1 with payload
{"blockIfFull":false,"_csrf_token":"36acbee2-6b78-4c13-b114-41040809833a"}
06:29:03,878 DEBUG HttpRequestExecutor - received HTTP status 202
06:29:03,878 DEBUG HttpRequestExecutor -
{"futureId":"898d546e-583f-4d37-9ca9-d1e10134037f"}
06:29:04,135 DEBUG PgxRemoteFuture - Requesting GET https://localhost:7007/
core/v1/futures/x-future-id/status HTTP/1.1
06:29:04,828 DEBUG PgxRemoteFuture - Requesting GET https://localhost:7007/
core/v1/futures/x-future-id/value HTTP/1.1
06:29:04,858 DEBUG RemoteUtils - received HTTP status 201
06:29:04,859 DEBUG RemoteUtils - {"id":"BE960B34-E135-4CF8-AB2F-
E1A6E2D7DB60","links":[{"href":"https://localhost:7007/core/v1/graphs/x-
graph-id","rel":"self","method":"GET","interaction":["async-polling"]},
{"href":"https://localhost:7007/core/v1/graphs/x-graph-
id","rel":"canonical","method":"GET","interaction":["async-
polling"]}],{"graphName":"bank_graph","vertexTables":{"Accounts":
{"name":"Accounts","metaData":{"name":"Accounts","idType":"integer","labels":
["Accounts"],"properties":[],"edgeProviderNamesWhereSource":
["Transfers"],"edgeProviderNamesWhereDestination":

```

```

["Transfers"],"id":null,"links":null},"providerLabels":
["Accounts"],"entityKeyType":"integer","isIdentityKeyMapping":false,"ve
rtexProperties":{},"vertexLabels":{"id":"19D95502-40D5-47F2-9F45-
B1CD09ECB989","links":[{"href":"https://localhost:7007/core/v1/
graphs/x-graph-id/properties/x-property-
name","rel":"self","method":"GET","interaction":["async-polling"]},
{"href":"https://localhost:7007/core/v1/graphs/x-graph-id/properties/x-
property-name","rel":"canonical","method":"GET","interaction":["async-
polling"]}], "dimension":-1,"propertyId":"19D95502-40D5-47F2-9F45-
B1CD09ECB989","name":"__vertex_labels__","entityType":"vertex","type":
"ro_string_set","namespace":"2C17C639-3771-3E30-88AE-34D6B380C5EC","tran
sient":false},"transient":false},"edgeTables":{"Transfers":
{"name":"Transfers","metaData":
{"name":"Transfers","idType":"long","directed":true,"labels":
["Transfers"],"properties":
[{"name":"AMOUNT","id":null,"propertyType":"float","dimension":0,"trans
ient":true,"links":null,"propertyId":"9A49BC0C-F8AA-465A-B8D6-
CA5A92BAE2C9"}]},"sourceVertexProviderName":"Accounts","destinationVerte
xProviderName":"Accounts","id":null,"links":null},"providerLabels":
["Transfers"],"entityKeyType":"long","isIdentityKeyMapping":true,"sourc
eVertexTableName":"Accounts","destinationVertexTableName":"Accounts","e
dgeProperties":{"FED6FE43-D311-46B6-9A5A-E8DC0D7B56C6":{"id":"FED6FE43-
D311-46B6-9A5A-E8DC0D7B56C6","links":[{"href":"https://localhost:7007/
core/v1/graphs/x-graph-id/properties/x-property-
name","rel":"self","method":"GET","interaction":["async-polling"]},
{"href":"https://localhost:7007/core/v1/graphs/x-graph-id/properties/x-
property-name","rel":"canonical","method":"GET","interaction":["async-
polling"]}], "dimension":0,"propertyId":"FED6FE43-D311-46B6-9A5A-
E8DC0D7B56C6","name":"AMOUNT","entityType":"edge","type":"float","names
pace":"2C17C639-3771-3E30-88AE-34D6B380C5EC","transient":false},"edgeL
abel":{"id":"371D2AC6-4EC5-45AD-8885-B3590F56D944","links
$5 ==> PgxGraph[name=bank_graph,N=1000,E=5001,created=1621160944599]

```

Python

```

>>>setloglevel("oracle.pgx","DEBUG")
>>>session.read_graph_with_properties("/scratch/PG/data/
bank_graph_analytics.json")
10:37:46.308 [main] DEBUG oracle.pgx.vfs.CommonsVfsProvider - resolve /
scratch/PG/data/bank_graph_analytics.json
10:37:46.314 [main] DEBUG oracle.pgx.config.AbstractConfigFactory -
parse graph config from /scratch/PG/data/bank_graph_analytics.json
(parent: file:///scratch/PG/data)
10:37:46.464 [main] DEBUG oracle.pgx.client.RemoteUtils - create
session cookie (session ID = 786242ef-a430-4964-8379-662079514d2e)
10:37:46.465 [main] DEBUG oracle.pgx.client.RemoteUtils - no value for
the sticky cookie given
10:37:46.466 [main] DEBUG oracle.pgx.client.RemoteUtils - create csrf
token cookie (token = f02ccd16-56e1-4e61-8d17-d6122f5ea48f)
10:37:46.565 [main] DEBUG oracle.pgx.client.HttpRequestExecutor -
Requesting POST https://localhost:7007/core/v1/loadGraph HTTP/1.1 with
payload
{"graphConfig":"TbPAV0ZGAB5yEUZcRkRQXERHDwJoTbTGU09tU1hVQFxaRghHfnwUHhZ
cBaTlQw4RcU5XVKNawUsRGxtGU09tRE5JUBMORlsLQ11RV0YQSURBD1VXWUNTGwxPD1tBU1

```

```

hZU2lZVU5mWFBEFFsLUBoMRkZHABsBQ1BSRExWWEVRckxSVVVIaVhTXVIbDxN2JXwuaH1lcXtnKzJ
+QxgRVkJGVFdADwIRRV1PW0UQHBJR15EFxBfbEMURk1CAEQXQ1ldREhTXEQWARpdVlRIFAwQeXmb
SB1PrkYcR10UCBZBERRED1MRHA9aWftRDwIRexhgxcRpbUpkGRNRAFUAaEhEXUJbAQnfEhYJa1YWU
lNNc1tcW0xAWBQIEl5dFx0WC1MIUhoMEGBABAheB1FBQw8YG11RVGdHTklIFAwQXfXUhmYRl4KVl
xfXFMQXx0PAkZWUVlRZl1RVGdeVklDX1hVEg1fVF1HAU9JFVxTQUBbCwdZCftdb1tRS0JRVWdDRVZ
bXlJXQhUDF3BxB10QWUxFEBgQAQNeFV1dUVldVlhrTldfQlRDFAwQZHHmdHJ3MG0scxoaEFBTEQdP
AEdWb1lVW1pRc1ZSWlWpDBRwcXlyamVsKmFHGxpFXUFABgNyAltFRUBaGwwWa2p8emZsdXVmb359F
x0WA10XWllCEA4QFwJPDEcRHA9ES1lEXhoJbEIPQk9CVRUDF1dYC1MRFRQUXfVfAEQXQ3V+f3h6bR
RJAUMRQ0BdUxQIEkRNR1haAxBJFVZXX1EQX0RpJGdwYmRkbX97YxpOahUPRVlHQ1RcakdRFkYAT2d
GQFtEDAJIEXYJEmxXWllBQ0xAFURwGhRcUvPcFwsWBlMLXGdRQFVCDTlMD1VfSVldWkUWUDgzNzmC
qx+4", "graphName": null, "_csrf_token": "f02ccd16-56e1-4e61-8d17-d6122f5ea48f"}
10:37:47.082 [main] DEBUG oracle.pgx.client.HttpRequestExecutor - received
HTTP status 202
10:37:47.082 [main] DEBUG oracle.pgx.client.HttpRequestExecutor -
{"futureId": "aeae66b7-e1d6-46f3-b78f-1c0d9642d308"}
10:37:47.086 [pgx-client-thread-2] DEBUG oracle.pgx.client.PgxRemoteFuture -
Requesting GET https://localhost:7007/core/v1/futures/x-future-id/status
HTTP/1.1
10:37:50.434 [pgx-client-thread-2] DEBUG oracle.pgx.client.PgxRemoteFuture -
Requesting GET https://localhost:7007/core/v1/futures/x-future-id/value
HTTP/1.1
10:37:50.539 [pgx-client-thread-2] DEBUG oracle.pgx.client.RemoteUtils -
received HTTP status 201
10:37:50.539 [pgx-client-thread-2] DEBUG oracle.pgx.client.RemoteUtils -
{"id": "2EE6F933-5679-4387-B79B-A4AAD0814DC6", "links": [{"href": "https://
localhost:7007/core/v1/graphs/x-graph-
id", "rel": "self", "method": "GET", "interaction": ["async-polling"]},
{"href": "https://localhost:7007/core/v1/graphs/x-graph-
id", "rel": "canonical", "method": "GET", "interaction": ["async-
polling"]}], "graphName": "bank_graph_analytics", "vertexTables": {"Accounts":
{"name": "Accounts", "metaData": {"name": "Accounts", "idType": "integer", "labels":
["Accounts"], "properties":
[{"name": "ID", "id": null, "propertyType": "integer", "dimension": 0, "transient": tr
ue, "links": null, "propertyId": "C933AB01-358B-4553-974D-0C54845719F2"},
{"name": "NAME", "id": null, "propertyType": "string", "dimension": 0, "transient": tr
ue, "links": null, "propertyId": "E0EDF68C-5AFE-4886-
B3A0-385CE56F1814"}], "edgeProviderNamesWhereSource":
["Transfers"], "edgeProviderNamesWhereDestination":
["Transfers"], "id": null, "links": null}, "providerLabels":
["Accounts"], "keyPropertyName": null, "entityKeyType": "integer", "isIdentityKeyM
apping": false, "vertexProperties": {"F5023C5A-5294-4383-BA4F-0109C67A020F":
{"id": "F5023C5A-5294-4383-BA4F-0109C67A020F", "links": [{"href": "https://
localhost:7007/core/v1/graphs/x-graph-id/properties/x-property-
name", "rel": "self", "method": "GET", "interaction": ["async-polling"]},
{"href": "https://localhost:7007/core/v1/graphs/x-graph-id/properties/x-
property-name", "rel": "canonical", "method": "GET", "interaction": ["async-
polling"]}], "dimension": 0, "propertyId": "F5023C5A-5294-4383-
BA4F-0109C67A020F", "name": "ID", "entityType": "vertex", "type": "integer", "namesp
ace": "2C17C639-3771-3E30-88AE-34D6B380C5EC", "transient": false}, {"C588F89E-53EB
-46DD-A83D-1078138C42C7": {"id": "C588F89E-53EB-46DD-
A83D-1078138C42C7", "links": [{"href": "https://localhost:7007/core/v1/graphs/x-
graph-id/properties/x-property-
name", "rel": "self", "method": "GET", "interaction": ["async-polling"]},
{"href": "https://localhost:7007/core/v1/graphs/x-graph-id/properties/x-
property-name", "rel": "canonical", "method": "GET", "interaction": ["async-

```

```

polling"]}], "dimension":0, "propertyId":"C588F89E-53EB-46DD-
A83D-1078138C42C7", "name":"NAME", "entityType":"vertex", "type":"string",
"namespace":"2C17C639-3771-3E30-88AE-34D6B380C5EC", "transient":false}},
"vertexLabels":{"id":"89FB1A38-BCDF-4FB0-9F8C-59471872BEA3", "links":
[{"href":"https://localhost:7007/core/v1/graphs/x-graph-id/
properties/x-property-name", "rel":"self", "method":"GET", "interaction":
["async-polling"]}, {"href":"https://localhost:7007/core/v1/graphs/x-
graph-id/properties/x-property-name", "rel":"canonical", "method":"GET
10:37:50.655 [pgx-client-thread-2] DEBUG oracle.pgx.api.PgxSession -
engine reports latest snapshot is 0 milli-seconds old. Max age is 0
milli-seconds
10:37:50.655 [pgx-client-thread-2] DEBUG oracle.pgx.api.PgxSession -
=> within range. Return snapshot
PgxGraph(name: bank_graph_analytics, v: 1000, e: 5001, directed: True,
memory(Mb): 0)

```

4.13.2 Connecting with Java

You can obtain a connection to a remote graph server (PGX) instance by simply passing the base URL of the remote PGX instance to the `getInstance()` method. By doing this, your application automatically uses the PGX client libraries to connect to a remotely-located graph server (PGX).

You can specify the base URL when you initialize the in-memory graph server (PGX) instance using Java. An example is as follows. A URL to an in-memory graph server (PGX) is provided to the `getInMemAnalyst` API call.

```

import oracle.pgx.api.*;
import oracle.pg.rdbms.*;
ServerInstance instance = GraphServer.getInstance("https://
<hostname>:<port>", "<username>", "<password>".toCharArray());
PgxSession session = instance.createSession("my-session");

```

- [Starting and Stopping the PGX Engine](#)

4.13.2.1 Starting and Stopping the PGX Engine

You can start the graph server (PGX) from the application by making a call to `instance.startEngine()` which takes a JSON object as an argument for PGX configuration.

Note:

- See [Connecting with Java](#) for more information about connecting to a graph server (PGX) instance and obtaining a `ServerInstance` object.
- See [Configuration Parameters for the Graph Server \(PGX\) Engine](#) for the various configuration options for the graph server (PGX).

Stopping the PGX Engine

You can stop the PGX engine using one of the following APIs:

```
instance.shutdownEngineNow(); // cancels pending tasks, throws exception if
engine is not running
instance.shutdownEngineNowIfRunning(); // cancels pending tasks, only tries
to shut down if engine is running
if (instance.shutdownEngine(30, TimeUnit.SECONDS) == false) {
    // doesn't accept new tasks but finishes up remaining tasks
    // pending tasks didn't finish after 30 seconds
}
```

Note:

Shutting down the PGX engine keeps the Apache Tomcat server alive, but new sessions cannot be created. Also, all the current sessions and tasks will be cancelled and terminated.

4.13.3 Connecting with Python

You can connect to a remote graph server (PGX) instance in your Python program. You must first authenticate with the remote server before you can create a session as illustrated in the following example:

```
import pypgx as pgx
import pypgx.pg.rdbms.graph_server as graph_server

base_url = "https://localhost:7007"
username = "scott"
password = "<password-for-scott>"

instance = graph_server.get_instance(base_url, username, password)
session = instance.create_session("python_pgx_client")
print(session)
```

To execute, save the above program into a file named `program.py` and run the following command:

```
python3 program.py
```

After successful login, you'll see the following message indicating a PGX session was created:

```
PgxSession(id: 0bdd4828-c3cc-4cef-92c8-0fcd105416f0, name: python_pgx_client)
```

**Note:**

To view the complete set of available Python APIs, see [Pypgx API](#).

4.14 Using Graph Server (PGX) as a Library

When you utilize PGX as a library in your application, the graph server (PGX) instance runs in the same JVM as the Java application and all requests are translated into direct function calls instead of remote procedure invocations.

In this case, you must install the graph server (PGX) using `RPM` in the same machine as the client applications. The shell executables provided by the graph server installation helps you to launch the Java or the Python shell in an embedded server mode. See [Installing Oracle Graph Server](#) for more information.

You can now start the Java shell without any parameters as shown:

```
cd /opt/oracle/graph
./bin/opg4j
```

The local PGX instance will try to load a PGX configuration file from:

```
/etc/oracle/graph/pgx.conf
```

You can change the location of the configuration file by passing the `--pgx_conf` command-line option followed by the path to the configuration file:

```
# start local PGX instance with custom config
./bin/opg4j --pgx_conf <path_to_pgx.conf>
```

You can also start the Python shell without any parameters as shown:

```
cd /opt/oracle/graph/
./bin/opg4py
```

When using Java, you can obtain a reference to the local PGX instance as shown:

```
import oracle.pg.rdbms.*;
import oracle.pgx.api.*;
...
ServerInstance instance = GraphServer.getEmbeddedInstance();
```

In a Python application, you can obtain a reference to the local PGX instance as shown:

```
import os
os.environ["PGX_CLASSPATH"] = "/opt/oracle/graph/lib/*"
import pypgx.pg.rdbms.graph_server as graph_server
...
instance = graph_server.get_embedded_instance()
```

Starting the PGX Engine

PGX provides a convenience mechanism to start the PGX Engine when using the graph server (PGX) as a library. That is, the graph server (PGX) is automatically initialized and starts up automatically when `ServerInstance.createSession()` is called the first time. This is provided that the engine is not already running at that time.

For this implicit initialization, PGX will configure itself with the PGX configuration file at the default locations. If the PGX configuration file is not found, PGX will configure itself using default parameter values as shown in [Configuration Parameters for the Graph Server \(PGX\) Engine](#).

Stopping the PGX Engine

When using the graph server (PGX) as a library, the `shutdownEngine()` method will be called automatically via a JVM shutdown hook on exit. Specifically, the shutdown hook is invoked once all the [non-daemon threads](#) of the application exit.

It is recommended that you do not terminate your PGX application forcibly with `kill -9`, as it will not clear the `temp` directory. See `tmp_dir` in [Configuration Parameters for the Graph Server \(PGX\) Engine](#).

4.15 User-Defined Functions (UDFs) in PGX

User-defined functions (UDFs) allow users of PGX to add custom logic to their PGQL queries or custom graph algorithms, to complement built-in functions with custom requirements.

▲ Caution:

UDFs enable running arbitrary code in the PGX server, possibly accessing sensitive data. Additionally, any PGX session can invoke any of the UDFs that are enabled on the PGX server. The application administrator who enables UDFs is responsible for checking the following:

- All the UDF code can be trusted.
- The UDFs are stored in a secure location that cannot be tampered with.

Furthermore, PGX assumes UDFs to be state-less and side-effect free.

PGX supports two types of UDFs:

- Java UDFs
- JavaScript UDFs

How to Use Java UDFs

The following simple example shows how to register a Java UDF at the PGX server and invoke it.

1. Create a class with a public static method. For example:

```
package my.udfs;

public class MyUdfs {
    public static String concat(String a, String b) {
        return a + b;
    }
}
```

2. Compile the class and compress into a JAR file. For example:

```
mkdir ./target
javac -d ./target *.java
cd target
jar cvf MyUdfs.jar *
```

3. Copy the JAR file into `/opt/oracle/graph/pgx/server/lib`.
4. Create a UDF JSON configuration file. For example, assume that `/path/to/my/udfs/dir/my_udfs.json` contains the following:

```
{
  "user_defined_functions": [
    {
      "namespace": "my",
      "language": "java",
      "implementation_reference": "my.udfs.MyUdfs",
      "function_name": "concat",
      "return_type": "string",
      "arguments": [
        {
          "name": "a",
          "type": "string"
        },
        {
          "name": "b",
          "type": "string"
        }
      ]
    }
  ]
}
```

5. Point to the directory containing the UDF configuration file in `/etc/oracle/graph/pgx.conf`. For example:

```
"udf_config_directory": "/path/to/my/udfs/dir/"
```

6. Restart the PGX server. For example:

```
sudo systemctl restart pgx
```

7. Try to invoke the UDF from within a PGQL query. For example:

```
graph.queryPgql("SELECT my.concat(my.concat(n.firstName, ' '),
n.lastName) FROM MATCH (n:Person)")
```

8. Try to invoke the UDF from within a PGX algorithm. For example:

 **Note:**

For each UDF you want to use, you need to create an abstract method with the same schema that gets annotated with the `@Udf` annotation.

```
import oracle.pgx.algorithm.annotations.Udf;
....

@GraphAlgorithm
public class MyAlgorithm {
    public void bomAlgorithm(PgxGraph g, VertexProperty<String> firstName,
VertexProperty<String> lastName, @Out VertexProperty<String> fullName) {

        ... fullName.set(v, concat(firstName.get(v), lastName.get(v))); ...

    }

    @Udf(namespace = "my")
    abstract String concat(String a, String b);
}
```

JavaScript UDFs

The requirements for a JavaScript UDF is as follows:

- The JavaScript source must contain all dependencies.
- The source must contain at least one valid export.
- The language parameter must be set to `javascript` in the UDF configuration file.

For example, consider a JavaScript source file `format.js` as shown:

```
//format.js
const fun = function(name, country) {
    if (country == null) return name;
    else return name + " (" + country + ")";
}

module.exports = {stringFormat: fun};
```

In order to load the UDF from `format.js`, the UDF configuration file will appear as follows:

```
{
  "namespace": "my",
```

```

"function_name": "format",
"language": "javascript",
"source_location": "format.js",
"source_function_name": "stringFormat",
"return_type": "string",
"arguments": [
  {
    "name": "name",
    "type": "string"
  },
  {
    "name": "country",
    "type": "string"
  }
]
}

```

 **Note:**

In this case, since the name of the UDF and the implementing method differ, you need to set the name of the UDF in the `source_function_name` field. Also, you can provide the path of the source code file in the `source_location` field.

UDF Configuration File Information

A UDF configuration file is a JSON file containing an array of `user_defined_functions`. (An example of such a file is in the step to "Create a UDF JSON configuration file" in the preceding [How to Use Java UDFs](#) subsection.)

Each user-defined function supports the fields shown in the following table.

Table 4-7 Fields for Each UDF

Field	Data Type	Description	Required?
<code>function_name</code>	string	Name of the function used as identifier in PGX	Required
<code>language</code>	enum[<code>java</code> , <code>javascript</code>]	Source language for the function (<code>java</code> or <code>javascript</code>)	Required
<code>return_type</code>	enum[<code>boolean</code> , <code>integer</code> , <code>long</code> , <code>float</code> , <code>double</code> , <code>string</code>]	Return type of the function	Required
<code>arguments</code>	array of object	Array of arguments. For each argument: type, argument name, required?	[]
<code>implementation_reference</code>	string	Reference to the function name on the classpath	null
<code>namespace</code>	string	Namespace of the function in PGX	null
<code>source_code</code>	string	Source code of the function provided inline	null

Table 4-7 (Cont.) Fields for Each UDF

Field	Data Type	Description	Required?
source_function_name	string	Name of the function in the source language	null
source_location	string	Local file path to the function's source code	null

All configured UDFs must be unique with regard to the combination of the following fields:

- namespace
- function_name
- arguments

4.16 Using HAProxy for PGX Load Balancing and High Availability

HAProxy is a high-performance TCP/HTTP load balancer and proxy server that allows multiplexing incoming requests across multiple web servers.

You can use HAProxy with multiple instances of the in-memory analytics server (PGX) for high availability. The following example uses the `opg` shell to connect to PGX.

The following instructions assume you have already installed and configured the in-memory analyst server, as explained in [Starting the In-Memory Graph Server \(PGX\)](#).

1. If HAProxy is not already installed on Big Data Appliance or your Oracle Linux distribution, run this command:

```
yum install haproxy
```

2. Start the PGX servers.
For example, if you want to load balance PGX across 4 nodes (such as `bda02`, `bda03`, `bda04`, and `bda05`) in the Big Data Appliance, start PGX on each of these nodes. Configure PGX to listen for connections on port 7007.
3. Configure HAProxy.
In this example, you will configure HAProxy to run on host `bda01` and to listen for incoming connections on port 8888. Create a new file `haproxy.cfg` on host `bda01` with the following content:

```
global
    maxconn 50000
    log /dev/log local0

defaults
    mode http
    option httplog
    log global
    option forwardfor
    timeout connect 5s
    timeout client 5s
```

```
timeout server 5s
balance source
hash-type consistent

listen www
  bind :8888
  server web1 bda02:7007 check
  server web2 bda03:7007 check
  server web3 bda04:7007 check
  server web4 bda05:7007 check
```

Specifying `balance source` maps the clients' IP addresses to corresponding servers' IP addresses. This is important because the PGX server relies on session stickiness during an analytics session. (For more information about configuring HAProxy, see the [HAProxy official documentation](#).)

4. Start the load balancer.
Start HAProxy on bda01 by passing in configuration file that you created in the preceding step:

```
haproxy -f haproxy.cfg
```

5. Test the load balancer.
From any host you can test connectivity to the HAProxy server by passing in the host and port of the server running HAProxy as the `base_url` parameter to the `opg` client shell. For example:

```
cd /opt/oracle/oracle-spatial-graph/property_graph
./bin/opg --base_url http://bda01:8888
```

 **Note:**

The PGX in-memory state is lost if the server goes down. HAProxy will route commands to another server, but the client must reload all graph data.

It is recommended that you run a series of PGX commands to test session affinity. Kill a server and restart the `opg` shell to confirm that HAProxy redirects the request to a new server.

5

Using the Property Graph Schema

This chapter provides conceptual and usage information about creating, storing, and working with property graph data in an Oracle Database environment.

You can create a property graph and store it in the property graph schema in Oracle Database in one of the following ways:

1. Use the `CREATE PROPERTY GRAPH` statement to create and populate these property graph schema objects.
2. Use `OPG_APIS.CREATE_PG`, to create the property graph schema objects. Then load data from the database tables into the schema objects using SQL or using the Data Access Layer APIs. The property graph schema provides a flexible schema option for storing your graph.

 **Note:**

The original database tables remain as-is and the data is copied from the original tables into the property graph schema tables.

- [Property Graph Schema Objects for Oracle Database](#)
The property graph PL/SQL and Java APIs use special Oracle Database schema objects.
- [Data Access Layer](#)
- [Getting Started with Property Graphs](#)
Follow these steps to get started with property graphs.
- [Using Java APIs for Property Graph Data](#)
Creating a property graph involves using the Java APIs to create the property graph and objects in it.
- [Managing Text Indexing for Property Graph Data](#)
Indexes in Oracle Spatial and Graph property graph support allow fast retrieval of elements by a particular key/value or key/text pair. These indexes are created based on an element type (vertices or edges), a set of keys (and values), and an index type.
- [Access Control for Property Graph Data \(Graph-Level and OLS\)](#)
Oracle Graph supports two access control and security models: graph level access control, and fine-grained security through integration with Oracle Label Security (OLS).
- [SQL-Based Property Graph Query and Analytics](#)
You can use SQL to query property graph data in Oracle Spatial and Graph.
- [Creating Property Graph Views on an RDF Graph](#)
With Oracle Graph, you can view RDF data as a property graph to execute graph analytics operations by creating property graph views over an RDF graph stored in Oracle Database.

5.1 Property Graph Schema Objects for Oracle Database

The property graph PL/SQL and Java APIs use special Oracle Database schema objects.

This topic describes objects related to the property graph schema approach to working with graph data.

Oracle Spatial and Graph lets you store, query, manipulate, and query property graph data in Oracle Database. For example, to create a property graph named `myGraph`, you can use either the Java APIs (`oracle.pg.rdbms.OraclePropertyGraph`) or the PL/SQL APIs (`MDSYS.OPG_APIS` package).

With the PL/SQL API:

```
BEGIN
    opg_apis.create_pg(
        'myGraph',
        dop => 4,                -- degree of parallelism
        num_hash_ptns => 8,     -- number of hash partitions used to
store the graph
        tbs => 'USERS',        -- tablespace
        options => 'COMPRESS=T'
    );
END;
/
```

With the Java API:

```
cfg = GraphConfigBuilder
    .forPropertyGraphRdbms()
    .setJdbcUrl("jdbc:oracle:thin:@127.0.0.1:1521:orcl")
    .setUsername("<your_user_name>")
    .setPassword("<your_password>")
    .setName("myGraph")
    .setMaxNumConnections(8)
    .setLoadEdgeLabel(false)
    .build();

OraclePropertyGraph opg = OraclePropertyGraph.getInstance(cfg);
```

- [Property Graph Tables \(Detailed Information\)](#)
- [Default Indexes on Vertex \(VT\\$\) and Edge \(GE\\$\) Tables](#)
- [Flexibility in the Property Graph Schema](#)

5.1.1 Property Graph Tables (Detailed Information)

After a property graph is established in the database, several tables are created automatically in the user's schema, with the graph name as the prefix and VT\$ or GE\$ as the suffix. For example, for a graph named `myGraph`, table `myGraphVT$` is created

to store vertices and their properties (K/V pairs), and table myGraphGE\$ is created to store edges and their properties.

Additional internal tables are created with IT\$ and GT\$ suffixes, to store text index metadata and graph skeleton (topological structure).

The definitions of tables myGraphVT\$ and myGraphGE\$ are as follows. They are important for SQL-based analytics and SQL-based property graph query. In both the VT\$ and GE\$ tables, VTS, VTE, and FE are reserved columns; column SL is for the security label; and columns K, T, V, VN, and VT together store all information about a property (K/V pair) of a graph element. In the VT\$ table, VID is a long integer for storing the vertex ID. In the GE\$ table, EID, SVID, and DVID are long integer columns for storing edge ID, source (from) vertex ID, and destination (to) vertex ID, respectively.

```
SQL> describe myGraphVT$
```

Name	Null?	Type
VID	NOT NULL	NUMBER
K		NVARCHAR2 (3100)
T		NUMBER (38)
V		NVARCHAR2 (15000)
VN		NUMBER
VT		TIMESTAMP (6) WITH TIME ZONE
SL		NUMBER
VTS		DATE
VTE		DATE
FE		NVARCHAR2 (4000)

```
SQL> describe myGraphGE$
```

Name	Null?	Type
EID	NOT NULL	NUMBER
SVID	NOT NULL	NUMBER
DVID	NOT NULL	NUMBER
EL		NVARCHAR2 (3100)
K		NVARCHAR2 (3100)
T		NUMBER (38)
V		NVARCHAR2 (15000)
VN		NUMBER
VT		TIMESTAMP (6) WITH TIME ZONE
SL		NUMBER
VTS		DATE
VTE		DATE
FE		NVARCHAR2 (4000)

For simplicity, only simple graph names are allowed, and they are case insensitive.

In both the VT\$ and GE\$ tables, Columns K, T, V, VN, VT together store all information about a property (K/V pair) of a graph element, while SL is used for security label, and VTS, VTE, FE are reserved columns.

In the property graph schema design, a property value is stored in the VN column if the value has numeric data type (long, int, double, float, and so on), in the VT column if the value is a

timestamp, or in the V column for Strings, boolean and other serializable data types. For better Oracle Text query support, a literal representation of the property value is saved in the V column even if the data type is numeric or timestamp. To differentiate all the supported data types, an integer ID is saved in the T column.

The K column in both VT\$ and GE\$ tables stores the property key. Each edge must have a label of String type, and the labels are stored in the EL column of the GE\$ table.

The T column in both VT\$ and GE\$ tables is a number representing the data type of the value of the property it describes. For example 1 means the value is a string, 2 means the value is an integer, and so on. Some T column possible values and associated data types are as follows:

- 1: STRING
- 2: INTEGER
- 3: FLOAT
- 4: DOUBLE
- 5: DATE
- 6: BOOLEAN
- 7: LONG
- 8: SHORT
- 9: BYTE
- 10: CHAR
- 20: Spatial data

The **VT\$ table** schema for storing vertices contains these columns:

- VID, a long column denoting the ID of the vertex.
- VL, a string column denoting the label of the vertex.
- K, a string column denoting the name of the property. If there is no property associated to the vertex, the value of this column should be a whitespace.
- T, a long column denoting the type of the property.
- V, a string column denoting the value of the property as a String. If the property type is numeric, a String format version of the value is stored in this column. Similarly, if the property is timestamp based, a String format version of the value is stored.
- VN, a numeric column denoting the value of a numeric property. This column stores the property value only if the property type is numeric.
- VT, a timestamp with time zone column storing the value of a date time property. This column stores the property value only if the property type is timestamp based.
- SL, a numeric column reserved for the security label set using Oracle Label Security (for further details on using Security Labels, see [Access Control for Property Graph Data \(Graph-Level and OLS\)](#)).
- VTS, a timestamp with time zone column reserved for future extensions.
- VTE, a timestamp with time zone column reserved for future extensions.
- FE, a string column reserved for future extensions.

The following example inserts rows into a table named CONNECTIONSVT\$. It includes T column values 1 through 10 (representing various data types).

```
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '1-STRING', 1,
'Some String', NULL, NULL);
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '2-INTEGER', 2,
NULL, 21, NULL);
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '3-FLOAT', 3,
NULL, 21.5, NULL);
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '4-DOUBLE', 4,
NULL, 21.5, NULL);
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '5-DATE', 5, NULL,
NULL, timestamp'2018-07-20 15:32:53.991000');
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '6-BOOLEAN', 6,
'Y', NULL, NULL);
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '7-LONG', 7, NULL,
42, NULL);
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '8-SHORT', 8,
NULL, 10, NULL);
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '9-BYTE', 9, NULL,
10, NULL);
INSERT INTO connectionsvt$(vid,k,t,v,vn,vt) VALUES (2001, '10-CHAR', 10,
'A', NULL, NULL);
...
UPDATE connectionsvt$ SET V = coalesce(v,to_nchar(vn),to_nchar(vt)) WHERE
vid=2001;
COMMIT;
```

The **GES** table schema for storing edges contains these columns:

- EID, a long column denoting the ID of the edge.
- SVID, a long column denoting the ID of the outgoing (origin) vertex.
- DVID, a long column denoting the ID of the incoming (destination) vertex.
- EL, a string column denoting the label of the edge.
- K, a string column denoting the name of the property. If there is no property associated to the vertex, the value of this column should be a whitespace.
- T, a long column denoting the type of the property.
- V, a string column denoting the value of the property as a String. If the property type is numeric, a String format version of the value is stored in this column. Similarly, if the property is timestamp based, a String format version of the value is stored.
- VN, a numeric column denoting the value of a numeric property. This column stores the property value only if the property type is numeric.
- VT, a timestamp with time zone column storing the value of a date time property. This column stores the property value only if the property type is timestamp based.
- SL, a numeric column reserved for the security label set using Oracle Label Security (for further details on using Security Labels, see [Access Control for Property Graph Data \(Graph-Level and OLS\)](#)).
- VTS, a timestamp with time zone column reserved for future extensions.
- VTE, a timestamp with time zone column reserved for future extensions.

- FE, a string column reserved for future extensions.

In addition to the VT\$ and GE\$ tables, Oracle Spatial and Graph maintains other internal tables.

An internal graph skeleton table, defined with the **GT\$ suffix**, is used to store the topological structure of a graph, and contains these columns:

- EID, a long column denoting the ID of the edge.
- EL, a string column denoting the label of the edge.
- SVID, a long column denoting the ID of the outgoing (origin) vertex.
- DVID, a long column denoting the ID of the incoming (destination) vertex.
- ELH, a raw column specifying the hash value of an edge label.
- ELS, a integer column specifying the edge label size with respect to total of characters.

An internal text index metadata table, created with **IT\$ suffix**, is used to store metadata information on text indexes created using the Oracle Text search engine. It is automatically populated based on the text indexes created. The IT\$ table includes the following columns for general information about a text index:

- EIN, a string column denoting the name of the text index.
- ET, a numeric column denoting the entities used to build the text index, if it is a vertex (1) or edge (2) text index.
- IT, a numeric column denoting the type of the text index, if it is an automatic (1) or manual (2) text index.
- SE, a numeric column denoting the search engine used to index the entities properties (2 indicates Oracle Text).
- K, a string column denoting the property name used for text indexing.

For Oracle Text-based indexes, the following columns are used to describe the configuration of the text index (for further details on building an Oracle Text-based index, see [Configuring Text Indexes Using Oracle Text](#)):

- PO, a column denoting the preferred owner for the text index configuration settings. By default, the package owner is set to MDSYS.
- DS, a string column specifying the data store used to build the text index.
- FIL, a string column specifying the filter used to build the text index.
- STR, a string column specifying the storage property used to build the text index.
- WL, a string column specifying the word list used when building the text index.
- SL, a string column specifying the stop list used to build the text index.
- LXR, a string column specifying the lexer used by Oracle Text during text indexing.
- OPTS, a string column specifying additional configuration options.

An internal table, defined with the **SS\$ suffix**, is created for Oracle internal use only.

5.1.2 Default Indexes on Vertex (VT\$) and Edge (GE\$) Tables

For query performance, several indexes on property graph tables are created by default. The index names follow the same convention as the table names, including

using the graph name as the prefix. For example, for the property graph `myGraph`, the following local (partitioned) indexes are created:

- A unique index `myGraphXQV$` on `myGraphVT$(VID, K)`
- A unique index `myGraphXQE$` on `myGraphGE$(EID, K)`
- An index `myGraphXSE$` on `myGraphGE$(SVID, DVID, EID, VN)`
- An index `myGraphXDE$` on `myGraphGE$(DVID, SVID, EID, VN)`

5.1.3 Flexibility in the Property Graph Schema

The property graph schema design does not use a catalog or centralized repository of any kind. Each property graph is separately stored and managed by a schema of user's choice. A user's schema may have one or more property graphs.

This design provides considerable flexibility to users. For example:

- Users can create additional indexes as desired.
- Different property graphs can have a different set of indexes or compression options for the base tables.
- Different property graphs can have different numbers of hash partitions.
- You can even drop the `XSE$` or `XDE$` index for a property graph; however, for integrity you should keep the unique constraints.

5.2 Data Access Layer

The data access layer provides a set of Java APIs that you can use to create and drop property graphs, add and remove vertices and edges, search for vertices and edges using key-value pairs, create text indexes, and perform other manipulations.

For more information, see:

- [Managing Text Indexing for Property Graph Data](#)
- [Using Java APIs for Property Graph Data](#)
- [Property Graph Schema Objects for Oracle Database \(PL/SQL and Java APIs\) and OPG_API\\$ Package Subprograms \(PL/SQL API\)](#).

5.3 Getting Started with Property Graphs

Follow these steps to get started with property graphs.

1. The first time you use property graphs, ensure that the software is installed and operational.
2. Interact with a graph using one or more of the following options:
 - Use Java APIs in your Java application. The Java APIs can also be run in the JShell Command line interface for prototype and demo purposes.
 - Run PGQL queries:
 - In the Java application, or
 - In the Graph visualization interface, or

- In the SQLcl client
- Run PGQL queries and execute Java APIs in the Apache Zeppelin interpreter
- [Required Privileges for Database Users](#)
The database schema that contains the graph tables (either Property Graph schema objects or relational tables that will be directly loaded as a graph in memory) requires certain privileges.

Related Topics

- [Using Java APIs for Property Graph Data](#)
Creating a property graph involves using the Java APIs to create the property graph and objects in it.

5.3.1 Required Privileges for Database Users

The database schema that contains the graph tables (either Property Graph schema objects or relational tables that will be directly loaded as a graph in memory) requires certain privileges.

```
ALTER SESSION  
CREATE PROCEDURE  
CREATE SEQUENCE  
CREATE SESSION  
CREATE TABLE  
CREATE TRIGGER  
CREATE TYPE  
CREATE VIEW
```

5.4 Using Java APIs for Property Graph Data

Creating a property graph involves using the Java APIs to create the property graph and objects in it.

- [Overview of the Java APIs](#)
- [Parallel Retrieval of Graph Data](#)
- [Using an Element Filter Callback for Subgraph Extraction](#)
- [Using Optimization Flags on Reads over Property Graph Data](#)
- [Adding and Removing Attributes of a Property Graph Subgraph](#)
- [Getting Property Graph Metadata](#)
- [Merging New Data into an Existing Property Graph](#)
- [Opening and Closing a Property Graph Instance](#)
- [Creating Vertices](#)
- [Creating Edges](#)
- [Deleting Vertices and Edges](#)
- [Reading a Graph from a Database into an Embedded In-Memory Analyst](#)
- [Specifying Labels for Vertices](#)
- [Building an In-Memory Graph](#)

- [Dropping a Property Graph](#)
- [Executing PGQL Queries](#)

5.4.1 Overview of the Java APIs

The Java APIs that you can use for property graphs include the following:

- [Oracle Graph Property Graph Java APIs](#)
- [Oracle Database Property Graph Java APIs](#)

5.4.1.1 Oracle Graph Property Graph Java APIs

Oracle Graph property graph support provides database-specific APIs for Oracle Database.

To use the Oracle Spatial and Graph API, import the following classes into your Java program:

```
import oracle.pg.common.*;
import oracle.pg.text.*;
import oracle.pg.rdbms.*;
import oracle.pg.rdbms.pgql.*;
import oracle.pgx.config.*;
import oracle.pgx.common.types.*;
```

To compile and run your Java applications, set your classpath to include the jar files in `<client-install-dir>/lib/`.

For example:

```
javac -cp ".:<client-install-dir>/lib/*" Main.java
java -cp ".:<client-install-dir>/lib/*" Main
```

5.4.1.2 Oracle Database Property Graph Java APIs

The Oracle Database property graph Java APIs enable you to create and populate a property graph stored in Oracle Database.

To use these Java APIs, import the classes into your Java program. For example:

```
import oracle.pg.rdbms.*;
import java.sql.*;
```

5.4.2 Parallel Retrieval of Graph Data

The parallel property graph query provides a simple Java API to perform parallel scans on vertices (or edges). Parallel retrieval is an optimized solution taking advantage of the distribution of the data across table partitions, so each partition is queried using a separate database connection.

Parallel retrieval will produce an array where each element holds all the vertices (or edges) from a specific partition (split). The subset of shards queried will be separated by the given start split ID and the size of the connections array provided. This way, the subset will consider splits in the range of `[start, start - 1 + size of connections array]`. Note that an integer ID (in the range of `[0, N - 1]`) is assigned to all the splits in the vertex table with `N` splits.

The following code loads a property graph, opens an array of connections, and executes a parallel query to retrieve all vertices and edges using the opened connections. The number of calls to the `getVerticesPartitioned` (`getEdgesPartitioned`) method is controlled by the total number of splits and the number of connections used.

```
OraclePropertyGraph opg = OraclePropertyGraph.getInstance(args,
szGraphName);

// Clear existing vertices/edges in the property graph
opg.clearRepository();

String szOPVFile = "../../data/connections.opv";
String szOPEFile = "../../data/connections.ope";

// This object will handle parallel data loading
OraclePropertyGraphDataLoader opgdl =
OraclePropertyGraphDataLoader.getInstance();
opgdl.loadData(opg, szOPVFile, szOPEFile, dop);

// Create connections used in parallel query
Oracle[] oracleConns = new Oracle[dop];
Connection[] conns = new Connection[dop];
for (int i = 0; i < dop; i++) {
    oracleConns[i] = opg.getOracle().clone();
    conns[i] = oracleConns[i].getConnection();
}

long lCountV = 0;
// Iterate over all the vertices' partitionIDs to count all the
vertices
for (int partitionID = 0; partitionID <
opg.getVertexPartitionsNumber();
    partitionID += dop) {
    Iterable<Vertex>[] iterables
        = opg.getVerticesPartitioned(conns /* Connection array */,
                                     true /* skip store to cache */,
                                     partitionID /* starting partition
*/);
    lCountV += consumeIterables(iterables); /* consume iterables using
threads */
}

// Count all vertices
System.out.println("Vertices found using parallel query: " + lCountV);

long lCountE = 0;
// Iterate over all the edges' partitionIDs to count all the edges
for (int partitionID = 0; partitionID <
opg.getEdgeTablePartitionIDs();
    partitionID += dop) {
    Iterable<Edge>[] iterables
        = opg.getEdgesPartitioned(conns /* Connection array */,
                                  true /* skip store to cache */,
                                  partitionID /* starting
```

```

partitionID */);
    lCountE += consumeIterables(iterables); /* consume iterables using
                                           threads */
}

// Count all edges
System.out.println("Edges found using parallel query: " + lCountE);

// Close the connections to the database after completed
for (int idx = 0; idx < conns.length; idx++) {
    conns[idx].close();
}

```

5.4.3 Using an Element Filter Callback for Subgraph Extraction

Oracle Spatial and Graph provides support for an easy subgraph extraction using user-defined element filter callbacks. An element filter callback defines a set of conditions that a vertex (or an edge) must meet in order to keep it in the subgraph. Users can define their own element filtering by implementing the `VertexFilterCallback` and `EdgeFilterCallback` API interfaces.

The following code fragment implements a `VertexFilterCallback` that validates if a vertex does not have a political role and its origin is the United States.

```

/**
 * VertexFilterCallback to retrieve a vertex from the United States
 * that does not have a political role
 */
private static class NonPoliticianFilterCallback
implements VertexFilterCallback
{
    @Override
    public boolean keepVertex(OracleVertexBase vertex)
    {
        String country = vertex.getProperty("country");
        String role = vertex.getProperty("role");

        if (country != null && country.equals("United States")) {
            if (role == null || !role.toLowerCase().contains("political")) {
                return true;
            }
        }

        return false;
    }

    public static NonPoliticianFilterCallback getInstance()
    {
        return new NonPoliticianFilterCallback();
    }
}

```

The following code fragment implements an `EdgeFilterCallback` that uses the `VertexFilterCallback` to keep only edges connected to the given input vertex, and whose connections are not politicians and come from the United States.

```

/**
 * EdgeFilterCallback to retrieve all edges connected to an input

```

```

    * vertex with "collaborates" label, and whose vertex is from the
    * United States with a role different than political
    */
private static class CollaboratorsFilterCallback
implements EdgeFilterCallback
{
private VertexFilterCallback m_vfc;
private Vertex m_startV;

public CollaboratorsFilterCallback(VertexFilterCallback vfc,
Vertex v)
{
m_vfc = vfc;
m_startV = v;
}

@Override
public boolean keepEdge(OracleEdgeBase edge)
{
if ("collaborates".equals(edge.getLabel())) {
if (edge.getVertex(Direction.IN).equals(m_startV) &&
m_vfc.keepVertex((OracleVertex)
edge.getVertex(Direction.OUT))) {
return true;
}
else if (edge.getVertex(Direction.OUT).equals(m_startV) &&
m_vfc.keepVertex((OracleVertex)
edge.getVertex(Direction.IN))) {
return true;
}
}
return false;
}

public static CollaboratorsFilterCallback
getInstance(VertexFilterCallback vfc, Vertex v)
{
return new CollaboratorsFilterCallback(vfc, v);
}
}

```

Using the filter callbacks previously defined, the following code fragment loads a property graph, creates an instance of the filter callbacks and later gets all of Robert Smith's collaborators who are not politicians and come from the United States.

```

OraclePropertyGraph opg = OraclePropertyGraph.getInstance(
args, szGraphName);

// Clear existing vertices/edges in the property graph
opg.clearRepository();

String szOPVFile = "../data/connections.opv";
String szOPEFile = "../data/connections.ope";

// This object will handle parallel data loading
OraclePropertyGraphDataLoader opgdL =
OraclePropertyGraphDataLoader.getInstance();
opgdL.loadData(opg, szOPVFile, szOPEFile, dop);

```

```

// VertexFilterCallback to retrieve all people from the United States // who are not
politicians
NonPoliticianFilterCallback npvfc = NonPoliticianFilterCallback.getInstance();

// Initial vertex: Robert Smith
Vertex v = opg.getVertices("name", "Robert Smith").iterator().next();

// EdgeFilterCallback to retrieve all collaborators of Robert Smith
// from the United States who are not politicians
CollaboratorsFilterCallback cefc = CollaboratorsFilterCallback.getInstance(npvfc, v);

Iterable<<Edge> smithCollabs = opg.getEdges((String[])null /* Match any
of the properties */,
cefc /* Match the
EdgeFilterCallback */
);
Iterator<<Edge> iter = smithCollabs.iterator();

System.out.println("\n\n-----Collaborators of Robert Smith from " +
" the US and non-politician\n\n");
long countV = 0;
while (iter.hasNext()) {
Edge edge = iter.next(); // get the edge
// check if smith is the IN vertex
if (edge.getVertex(Direction.IN).equals(v)) {
System.out.println(edge.getVertex(Direction.OUT) + "(Edge ID: " +
edge.getId() + ")"); // get out vertex
}
else {
System.out.println(edge.getVertex(Direction.IN) + "(Edge ID: " +
edge.getId() + ")"); // get in vertex
}

countV++;
}

```

By default, all reading operations such as get all vertices, get all edges (and parallel approaches) will use the filter callbacks associated with the property graph using the methods `opg.setVertexFilterCallback(vfc)` and `opg.setEdgeFilterCallback(efc)`. If there is no filter callback set, then all the vertices (or edges) and edges will be retrieved.

The following code fragment uses the default edge filter callback set on the property graph to retrieve the edges.

```

// VertexFilterCallback to retrieve all people from the United States // who are not
politicians
NonPoliticianFilterCallback npvfc = NonPoliticianFilterCallback.getInstance();

// Initial vertex: Robert Smith
Vertex v = opg.getVertices("name", "Robert Smith").iterator().next();

// EdgeFilterCallback to retrieve all collaborators of Robert Smith
// from the United States who are not politicians
CollaboratorsFilterCallback cefc = CollaboratorsFilterCallback.getInstance(npvfc, v);

opg.setEdgeFilterCallback(cefc);

Iterable<Edge> smithCollabs = opg.getEdges();
Iterator<Edge> iter = smithCollabs.iterator();

System.out.println("\n\n-----Collaborators of Robert Smith from " +

```

```

    " the US and non-politician\n\n");
    long countV = 0;
    while (iter.hasNext()) {
    Edge edge = iter.next(); // get the edge
    // check if smith is the IN vertex
    if (edge.getVertex(Direction.IN).equals(v)) {
        System.out.println(edge.getVertex(Direction.OUT) + "(Edge ID: " +
            edge.getId() + ")"); // get out vertex
    }
    else {
        System.out.println(edge.getVertex(Direction.IN) + "(Edge ID: " +
            edge.getId() + ")"); // get in vertex
    }

    countV++;
    }

```

5.4.4 Using Optimization Flags on Reads over Property Graph Data

Oracle Spatial and Graph provides support for optimization flags to improve graph iteration performance. Optimization flags allow processing vertices (or edges) as objects with none or minimal information, such as ID, label, and/or incoming/outgoing vertices. This way, the time required to process each vertex (or edge) during iteration is reduced.

The following table shows the optimization flags available when processing vertices (or edges) in a property graph.

Optimization Flag	Description
DO_NOT_CREATE_OBJECT	Use a predefined constant object when processing vertices or edges.
JUST_EDGE_ID	Construct edge objects with ID only when processing edges.
JUST_LABEL_EDGE_ID	Construct edge objects with ID and label only when processing edges.
JUST_LABEL_VERTEX_EDGE_ID	Construct edge objects with ID, label, and in/out vertex IDs only when processing edges
JUST_VERTEX_EDGE_ID	Construct edge objects with just ID and in/out vertex IDs when processing edges.
JUST_VERTEX_ID	Construct vertex objects with ID only when processing vertices.

The following code fragment uses a set of optimization flags to retrieve only all the IDs from the vertices and edges in the property graph. The objects retrieved by reading all vertices and edges will include only the IDs and no Key/Value properties or additional information.

```

import oracle.pg.common.OraclePropertyGraphBase.OptimizationFlag;
OraclePropertyGraph opg = OraclePropertyGraph.getInstance(
    args, szGraphName);

// Clear existing vertices/edges in the property graph
opg.clearRepository();

String szOPVFile = "../data/connections.opv";
String szOPEFile = "../data/connections.ope";

// This object will handle parallel data loading
OraclePropertyGraphDataLoader opgdl =

```

```

OraclePropertyGraphDataLoader.getInstance();
opgdl.loadData(opg, szOPVFile, szOPEFile, dop);

// Optimization flag to retrieve only vertices IDs
OptimizationFlag optFlagVertex = OptimizationFlag.JUST_VERTEX_ID;

// Optimization flag to retrieve only edges IDs
OptimizationFlag optFlagEdge = OptimizationFlag.JUST_EDGE_ID;

// Print all vertices
Iterator<Vertex> vertices =
opg.getVertices((String[])null /* Match any of the
properties */,
null /* Match the VertexFilterCallback */,
optFlagVertex /* optimization flag */
).iterator();

System.out.println("----- Vertices IDs-----");
long vCount = 0;
while (vertices.hasNext()) {
OracleVertex v = vertices.next();
System.out.println((Long) v.getId());
vCount++;
}
System.out.println("Vertices found: " + vCount);

// Print all edges
Iterator<Edge> edges =
opg.getEdges((String[])null /* Match any of the properties */,
null /* Match the EdgeFilterCallback */,
optFlagEdge /* optimization flag */
).iterator();

System.out.println("----- Edges -----");
long eCount = 0;
while (edges.hasNext()) {
Edge e = edges.next();
System.out.println((Long) e.getId());
eCount++;
}
System.out.println("Edges found: " + eCount);

```

By default, all reading operations such as get all vertices, get all edges (and parallel approaches) will use the optimization flag associated with the property graph using the method `opg.setDefaultVertexOptFlag(optFlagVertex)` and `opg.setDefaultEdgeOptFlag(optFlagEdge)`. If the optimization flags for processing vertices and edges are not defined, then all the information about the vertices and edges will be retrieved.

The following code fragment uses the default optimization flags set on the property graph to retrieve only all the IDs from its vertices and edges.

```

import oracle.pg.common.OraclePropertyGraphBase.OptimizationFlag;

// Optimization flag to retrieve only vertices IDs
OptimizationFlag optFlagVertex = OptimizationFlag.JUST_VERTEX_ID;

// Optimization flag to retrieve only edges IDs
OptimizationFlag optFlagEdge = OptimizationFlag.JUST_EDGE_ID;

```

```

opg.setDefaultVertexOptFlag(optFlagVertex);
opg.setDefaultEdgeOptFlag(optFlagEdge);

Iterator<Vertex> vertices = opg.getVertices().iterator();
System.out.println("----- Vertices IDs-----");
long vCount = 0;
while (vertices.hasNext()) {
    OracleVertex v = vertices.next();
    System.out.println((Long) v.getId());
    vCount++;
}
System.out.println("Vertices found: " + vCount);

// Print all edges
Iterator<Edge> edges = opg.getEdges().iterator();
System.out.println("----- Edges -----");
long eCount = 0;
while (edges.hasNext()) {
    Edge e = edges.next();
    System.out.println((Long) e.getId());
    eCount++;
}
System.out.println("Edges found: " + eCount);

```

5.4.5 Adding and Removing Attributes of a Property Graph Subgraph

Oracle Spatial and Graph supports updating attributes (key/value pairs) to a subgraph of vertices and/or edges by using a user-customized operation callback. An operation callback defines a set of conditions that a vertex (or an edge) must meet in order to update it (either add or remove the given attribute and value).

You can define your own attribute operations by implementing the `VertexOpCallback` and `EdgeOpCallback` API interfaces. You must override the `needOp` method, which defines the conditions to be satisfied by the vertices (or edges) to be included in the update operation, as well as the `getAttributeKeyName` and `getAttributeKeyValue` methods, which return the key name and value, respectively, to be used when updating the elements.

The following code fragment implements a `VertexOpCallback` that operates over the `smithCollaborator` attribute associated only with Robert Smith collaborators. The value of this property is specified based on the role of the collaborators.

```

private static class CollaboratorsVertexOpCallback
implements VertexOpCallback
{
    private OracleVertexBase m_smith;
    private List<Vertex> m_smithCollaborators;

    public CollaboratorsVertexOpCallback(OraclePropertyGraph opg)
    {
        // Get a list of Robert Smith's Collaborators
        m_smith = (OracleVertexBase) opg.getVertices("name",
            "Robert Smith")
            .iterator().next();

        Iterable<Vertex> iter = m_smith.getVertices(Direction.BOTH,
            "collaborates");
    }
}

```



```
m_smithCollaborators = OraclePropertyGraphUtils.listify(iter);
}

public static CollaboratorsVertexOpCallback
getInstance(OraclePropertyGraph opg)
{
return new CollaboratorsVertexOpCallback(opg);
}

/**
 * Add attribute if and only if the vertex is a collaborator of Robert
 * Smith
 */
@Override
public boolean needOp(OracleVertexBase v)
{
return m_smithCollaborators != null &&
    m_smithCollaborators.contains(v);
}

@Override
public String getAttributeKeyName(OracleVertexBase v)
{
return "smithCollaborator";
}

/**
 * Define the property's value based on the vertex role
 */
@Override
public Object getAttributeKeyValue(OracleVertexBase v)
{
String role = v.getProperty("role");
role = role.toLowerCase();
if (role.contains("political")) {
return "political";
}
else if (role.contains("actor") || role.contains("singer") ||
    role.contains("actress") || role.contains("writer") ||
    role.contains("producer") || role.contains("director")) {
return "arts";
}
else if (role.contains("player")) {
return "sports";
}
else if (role.contains("journalist")) {
return "journalism";
}
else if (role.contains("business") || role.contains("economist")) {
return "business";
}
else if (role.contains("philanthropist")) {
return "philanthropy";
}
return " ";
}
}
```

The following code fragment implements an `EdgeOpCallback` that operates over the `smithFeud` attribute associated only with Robert Smith feuds. The value of this property is specified based on the role of the collaborators.

```
private static class FeudsEdgeOpCallback
implements EdgeOpCallback
{
    private OracleVertexBase m_smith;
    private List<Edge> m_smithFeuds;

    public FeudsEdgeOpCallback(OraclePropertyGraph opg)
    {
        // Get a list of Robert Smith's feuds
        m_smith = (OracleVertexBase) opg.getVertices("name",
            "Robert Smith")
            .iterator().next();

        Iterable<Vertex> iter = m_smith.getVertices(Direction.BOTH,
            "feuds");
        m_smithFeuds = OraclePropertyGraphUtils.listify(iter);
    }

    public static FeudsEdgeOpCallback getInstance(OraclePropertyGraph opg)
    {
        return new FeudsEdgeOpCallback(opg);
    }

    /**
     * Add attribute if and only if the edge is in the list of Robert Smith's
     * feuds
     */
    @Override
    public boolean needOp(OracleEdgeBase e)
    {
        return m_smithFeuds != null && m_smithFeuds.contains(e);
    }

    @Override
    public String getAttributeKeyName(OracleEdgeBase e)
    {
        return "smithFeud";
    }

    /**
     * Define the property's value based on the in/out vertex role
     */
    @Override
    public Object getAttributeKeyValue(OracleEdgeBase e)
    {
        OracleVertexBase v = (OracleVertexBase) e.getVertex(Direction.IN);
        if (m_smith.equals(v)) {
            v = (OracleVertexBase) e.getVertex(Direction.OUT);
        }
        String role = v.getProperty("role");
        role = role.toLowerCase();

        if (role.contains("political")) {
            return "political";
        }
        else if (role.contains("actor") || role.contains("singer") ||
            role.contains("actress") || role.contains("writer") ||
            role.contains("producer") || role.contains("director")) {
            return "arts";
        }
        else if (role.contains("journalist")) {
```

```

return "journalism";
}
else if (role.contains("player")) {
return "sports";
}
else if (role.contains("business") || role.contains("economist")) {
return "business";
}
else if (role.contains("philanthropist")) {
return "philanthropy";
}
return " ";
}
}

```

Using the operations callbacks defined previously, the following code fragment loads a property graph, creates an instance of the operation callbacks, and later adds the attributes into the pertinent vertices and edges using the `addAttributeToAllVertices` and `addAttributeToAllEdges` methods in `OraclePropertyGraph`.

```

OraclePropertyGraph opg = OraclePropertyGraph.getInstance(
    args, szGraphName);

// Clear existing vertices/edges in the property graph
opg.clearRepository();

String szOPVFile = "../../data/connections.opv";
String szOPEFile = "../../data/connections.ope";

// This object will handle parallel data loading
OraclePropertyGraphDataLoader opgdL = OraclePropertyGraphDataLoader.getInstance();
opgdL.loadData(opg, szOPVFile, szOPEFile, dop);

// Create the vertex operation callback
CollaboratorsVertexOpCallback cvoc = CollaboratorsVertexOpCallback.getInstance(opg);

// Add attribute to all people collaborating with Smith based on their role
opg.addAttributeToAllVertices(cvoc, true /** Skip store to Cache */, dop);

// Look up for all collaborators of Smith
Iterable<Vertex> collaborators = opg.getVertices("smithCollaborator", "political");
System.out.println("Political collaborators of Robert Smith " +
    getVerticesAsString(collaborators));

collaborators = opg.getVertices("smithCollaborator", "business");
System.out.println("Business collaborators of Robert Smith " +
    getVerticesAsString(collaborators));

// Add an attribute to all people having a feud with Robert Smith to set
// the type of relation they have
FeudsEdgeOpCallback feoc = FeudsEdgeOpCallback.getInstance(opg);
opg.addAttributeToAllEdges(feoc, true /** Skip store to Cache */, dop);

// Look up for all feuds of Smith
Iterable<Edge> feuds = opg.getEdges("smithFeud", "political");
System.out.println("\n\nPolitical feuds of Robert Smith " + getEdgesAsString(feuds));

feuds = opg.getEdges("smithFeud", "business");
System.out.println("Business feuds of Robert Smith " +
    getEdgesAsString(feuds));

```

The following code fragment defines an implementation of `VertexOpCallback` that can be used to remove vertices having value `philanthropy` for attribute `smithCollaborator`, then call the API `removeAttributeFromAllVertices`; It also defines an implementation of `EdgeOpCallback` that can be used to remove edges having value `business` for attribute `smithFeud`, then call the API `removeAttributeFromAllEdges`.

```
System.out.println("\n\nRemove 'smithCollaborator' property from all the" +
    "philanthropy collaborators");
PhilanthropyCollaboratorsVertexOpCallback pvoc =
    PhilanthropyCollaboratorsVertexOpCallback.getInstance();

opg.removeAttributeFromAllVertices(pvoc);

System.out.println("\n\nRemove 'smithFeud' property from all the" + "business
feuds");
BusinessFeudsEdgeOpCallback beoc = BusinessFeudsEdgeOpCallback.getInstance();

opg.removeAttributeFromAllEdges(beoc);

/**
 * Implementation of a EdgeOpCallback to remove the "smithCollaborators"
 * property from all people collaborating with Robert Smith that have a
 * philanthropy role
 */
private static class PhilanthropyCollaboratorsVertexOpCallback implements
VertexOpCallback
{
    public static PhilanthropyCollaboratorsVertexOpCallback getInstance()
    {
        return new PhilanthropyCollaboratorsVertexOpCallback();
    }

    /**
     * Remove attribute if and only if the property value for
     * smithCollaborator is Philanthropy
     */
    @Override
    public boolean needOp(OracleVertexBase v)
    {
        String type = v.getProperty("smithCollaborator");
        return type != null && type.equals("philanthropy");
    }

    @Override
    public String getAttributeKeyName(OracleVertexBase v)
    {
        return "smithCollaborator";
    }

    /**
     * Define the property's value. In this case can be empty
     */
    @Override
    public Object getAttributeKeyValue(OracleVertexBase v)
    {
        return " ";
    }
}

/**
 * Implementation of a EdgeOpCallback to remove the "smithFeud" property
```

```

    * from all connections in a feud with Robert Smith that have a business role
    */
private static class BusinessFeudsEdgeOpCallback implements EdgeOpCallback
{
    public static BusinessFeudsEdgeOpCallback getInstance()
    {
        return new BusinessFeudsEdgeOpCallback();
    }

    /**
     * Remove attribute if and only if the property value for smithFeud is
     * business
     */
    @Override
    public boolean needOp(OracleEdgeBase e)
    {
        String type = e.getProperty("smithFeud");
        return type != null && type.equals("business");
    }

    @Override
    public String getAttributeKeyName(OracleEdgeBase e)
    {
        return "smithFeud";
    }

    /**
     * Define the property's value. In this case can be empty
     */
    @Override
    public Object getAttributeKeyValue(OracleEdgeBase e)
    {
        return " ";
    }
}

```

5.4.6 Getting Property Graph Metadata

You can get graph metadata and statistics, such as all graph names in the database; for each graph, getting the minimum/maximum vertex ID, the minimum/maximum edge ID, vertex property names, edge property names, number of splits in graph vertex, and the edge table that supports parallel table scans.

The following code fragment gets the metadata and statistics of the existing property graphs stored in an Oracle database.

```

// Get all graph names in the database
List<String> graphNames = OraclePropertyGraphUtils.getGraphNames(dbArgs);

for (String graphName : graphNames) {
    OraclePropertyGraph opg = OraclePropertyGraph.getInstance(args,
graphName);

    System.err.println("\n Graph name: " + graphName);
    System.err.println(" Total vertices: " +
        opg.countVertices(dop));

    System.err.println(" Minimum Vertex ID: " +
        opg.getMinVertexID(dop));
    System.err.println(" Maximum Vertex ID: " +

```

```

    opg.getMaxVertexID(dop));

Set<String> propertyNamesV = new HashSet<String>();
opg.getVertexPropertyNames(dop, 0 /* timeout,0 no timeout */,
    propertyNamesV);

System.err.println(" Vertices property names: " +
    getPropertyNamesAsString(propertyNamesV));

System.err.println("\n\n Total edges: " + opg.countEdges(dop));
System.err.println(" Minimum Edge ID: " + opg.getMinEdgeID(dop));
System.err.println(" Maximum Edge ID: " + opg.getMaxEdgeID(dop));

Set<String> propertyNamesE = new HashSet<String>();
opg.getEdgePropertyNames(dop, 0 /* timeout,0 no timeout */,
    propertyNamesE);

System.err.println(" Edge property names: " +
    getPropertyNamesAsString(propertyNamesE));

System.err.println("\n\n Table Information: ");
System.err.println("Vertex table number of splits: " +
    (opg.getVertexPartitionsNumber()));
System.err.println("Edge table number of splits: " +
    (opg.getEdgePartitionsNumber()));
}

```

5.4.7 Merging New Data into an Existing Property Graph

In addition to loading graph data into an empty property graph in Oracle Database, you can merge new graph data into an existing (empty or non-empty) graph. As with data loading, data merging splits the input vertices and edges into multiple chunks and merges them with the existing graph in database in parallel.

When doing the merging, the flows are different depends on whether there is an overlap between new graph data and existing graph data. *Overlap* here means that the same key of a graph element may have different values in the new and existing graph data. For example, key `weight` of the vertex with ID 1 may have value 0.8 in the new graph data and value 0.5 in the existing graph data. In this case, you must specify whether the new value or the existing value should be used for the key.

The following options are available for graph data merging: JDB-based, external table-based, and SQL loader-based merging.

- JDB-based Graph Data Merging
- External Table-Based Data Merging
- SQL Loader-Based Data Merging

JDB-based Graph Data Merging

JDB-based data merging uses Java Database Connectivity (JDB) APIs to load the new graph data into Oracle Database and then merge the new graph data into an existing graph.

The following example merges the new graph data from vertex and edge files `szOPVFile` and `szOPEFile` in Oracle-defined Flat-file format with an existing graph

named `opg`, using a JDBC-based data merging with a DOP (degree of parallelism) of 48, batch size of 1000, and specified data merging options.

```
String szOPVFile = "../../data/connectionsNew.opv";
String szOPEFile = "../../data/connectionsNew.ope";
OraclePropertyGraphDataLoader opgdl =
OraclePropertyGraphDataLoader.getInstance();
opgdl.mergeData(opg, szOPVFile, szOPEFile,
    48 /*DOP*/,
    1000 /*Batch Size*/,
    true /*Rebuild index*/,
    "pdml=t, pddl=t, no_dup=t, use_new_val_for_dup_key=t" /*Merge
options*/);
```

To optimize the performance of the data merging operations, a set of flags and hints can be specified in the merging options parameter when calling the JDBC-based data merging. These hints include:

- **DOP:** The degree of parallelism to use when merging the data. This parameter determines the number of chunks to generate when splitting the file, as well as the number of loader threads to use when merging the data into the property graph VT\$ and GE\$ tables.
- **Batch Size:** An integer specifying the batch size to use for Oracle JDBC statements in batching mode.
- **Rebuild index:** If set to true, the data loader will disable all the indexes and constraints defined over the property graph into which the data will be loaded. After all the data is merged into the property graph, all the original indexes and constraints will be rebuilt and enabled.
- **Merge options:** An option (or multiple options separated by commas) to optimize the data merging operations. These options include:
 - **PDML=T:** enables parallel execution for DML operations for the database session used in the data loader. This hint is used to improve the performance of long-running batching jobs.
 - **PDDL=T:** enables parallel execution for DDL operations for the database session used in the data loader. This hint is used to improve the performance of long-running batching jobs.
 - **NO_DUP=T:** assumes the input new graph data does not have invalid duplicates. In a valid property graph, each vertex (or edge) can at most have one value for a given property key. In an invalid property graph, a vertex (or edge) may have two or more values for a particular key. As an example, a vertex, *v*, has two key/value pairs: `name/"John"` and `name/"Johnny"`, and they share the same key.
 - **OVERLAP=F:** assumes there is no overlap between new graph data and existing graph data. That is, there is no key with multiple distinct values in the new and existing graph data.
 - **USE_NEW_VAL_FOR_DUP_KEY=T:** if there is overlap between new graph data and existing graph data, use the value in the new graph data; otherwise, use the value in the existing graph data.

External Table-Based Data Merging

External table-based data merging uses an external table to load new graph data into Oracle Database and then merge the new graph data into an existing graph.

External-table based data merging requires a directory object, where the files read by the external tables will be stored. This directory can be created using the following SQL*Plus statements:

```
create or replace directory tmp_dir as '/tmppath/';
grant read, write on directory tmp_dir to public;
```

The following example merges the new graph data from a vertex and edge files szOPVFile and szOPEFile in Oracle flat-file format with an existing graph opg using an external table-based data merging, a DOP (degree of parallelism) of 48, and specified merging options.

```
String szOPVFile = "../data/connectionsNew.opv";
String szOPEFile = "../data/connectionsNew.ope";
String szExtDir = "tmp_dir";
OraclePropertyGraphDataLoader opgdl =
OraclePropertyGraphDataLoader.getInstance();
opgdl.mergeDataWithExtTab(opg, szOPVFile, szOPEFile,
    48 /*DOP*/,
    true /*Use Named Pipe for splitting*/,
    szExtDir /*database directory object*/,
    true /*Rebuild index*/,
    "pdml=t, pddl=t, no_dup=t, use_new_val_for_dup_key=t" /*Merge
options*/);
```

SQL Loader-Based Data Merging

SQL loader-based data merging uses Oracle SQL*Loader to load the new graph data into Oracle Database and then merge the new graph data into an existing graph.

The following example merges the new graph data from a vertex and edge files szOPVFile and szOPEFile in Oracle Flat-file format with an existing graph opg using an SQL loader -based data merging with a DOP (degree of parallelism) of 48 and the specified merging options. To use the APIs, the path to the SQL*Loader needs to be specified.

```
String szUser = "username";
String szPassword = "password";
String szDbId = "db18c"; /*service name of the database*/
String szOPVFile = "../data/connectionsNew.opv"; 0
String szOPEFile = "../data/connectionsNew.ope";
String szSQLLoaderPath = "<YOUR_ORACLE_HOME>/bin/sqlldr";
OraclePropertyGraphDataLoader opgdl =
OraclePropertyGraphDataLoader.getInstance();
opgdl.mergeDataWithSqlLdr(opg, szUser, szPassword, szDbId, szOPVFile,
szOPEFile,
    48 /*DOP*/,
    true /*Use Named Pipe for splitting*/,
    szSQLLoaderPath /* SQL*Loader path: the path to bin/sqlldr */,
```



```

    true /*Rebuild index*/,
    "pdml=t, pddl=t, no_dup=t, use_new_val_for_dup_key=t" /*Merge
options*/);

```

5.4.8 Opening and Closing a Property Graph Instance

When describing a property graph, use these Oracle Property Graph classes to open and close the property graph instance properly:

- `OraclePropertyGraph.getInstance`: Opens an instance of an Oracle property graph. This method has two parameters, the connection information and the graph name. The format of the connection information depends on whether you use HBase or Oracle NoSQL Database as the backend database.
- `OraclePropertyGraph.clearRepository`: Removes all vertices and edges from the property graph instance.
- `OraclePropertyGraph.shutdown`: Closes the graph instance.

For Oracle Database, the `OraclePropertyGraph.getInstance` method uses an Oracle instance to manage the database connection. `OraclePropertyGraph` has a set of constructors that let you set the graph name, number of hash partitions, degree of parallelism, tablespace, and options for storage (such as compression). For example:

```

import oracle.pg.rdbms.*;
Oracle oracle = new Oracle(jdbcURL, username, password);

OraclePropertyGraph opg = OraclePropertyGraph.getInstance(oracle, graphName);
opg.clearRepository();
//      .
//      . Graph description
//      .
// Close the graph instance
opg.shutdown();

```

If the in-memory analyst functions are required for an application, you should use `GraphConfigBuilder` to create a graph for Oracle Database, and instantiate `OraclePropertyGraph` with that graph name as an argument. For example, the following code snippet constructs a graph `config`, gets an `OraclePropertyGraph` instance, loads some data into that graph, and gets an in-memory analyst.

```

import oracle.pgx.config.*;
import oracle.pgx.api.*;
import oracle.pgx.common.types.*;

...

PgNosqlGraphConfig cfg = GraphConfigBuilder.forPropertyGraphRdbms ()
    .setJdbcUrl("jdbc:oracle:thin:@<hostname>:1521:<sid>")
    .setUsername("<username>").setPassword("<password>")
    .setName(szGraphName)
    .setMaxNumConnections(8)
    .addEdgeProperty("lbl", PropertyType.STRING, "lbl")
    .addEdgeProperty("weight", PropertyType.DOUBLE, "1000000")
    .build();

```

```

OraclePropertyGraph opg = OraclePropertyGraph.getInstance(cfg);

String szOPVFile = "../data/connections.opv";
String szOPEFile = "../data/connections.ope";

// perform a parallel data load
OraclePropertyGraphDataLoader opgdl =
OraclePropertyGraphDataLoader.getInstance();
opgdl.loadData(opg, szOPVFile, szOPEFile, 2 /* dop */, 1000, true,
"PDML=T,PDDL=T,NO_DUP=T,");

...
PgxSession session = Pgx.createSession("session-id-1");
PgxGraph g = session.readGraphWithProperties(cfg);

Analyst analyst = session.createAnalyst();
...

```

5.4.9 Creating Vertices

To create a vertex, use these Oracle Property Graph methods:

- `OraclePropertyGraph.addVertex`: Adds a vertex instance to a graph.
- `OracleVertex.setProperty`: Assigns a key-value property to a vertex.
- `OraclePropertyGraph.commit`: Saves all changes to the property graph instance.

The following code fragment creates two vertices named `v1` and `v2`, with properties for age, name, weight, height, and sex in the `opg` property graph instance. The `v1` properties set the data types explicitly.

```

// Create vertex v1 and assign it properties as key-value pairs
Vertex v1 = opg.addVertex(11);
v1.setProperty("age", Integer.valueOf(31));
v1.setProperty("name", "Alice");
v1.setProperty("weight", Float.valueOf(135.0f));
v1.setProperty("height", Double.valueOf(64.5d));
v1.setProperty("female", Boolean.TRUE);

Vertex v2 = opg.addVertex(21);
v2.setProperty("age", 27);
v2.setProperty("name", "Bob");
v2.setProperty("weight", Float.valueOf(156.0f));
v2.setProperty("height", Double.valueOf(69.5d));
v2.setProperty("female", Boolean.FALSE);

```

5.4.10 Creating Edges

To create an edge, use these Oracle Property Graph methods:

- `OraclePropertyGraph.addEdge`: Adds an edge instance to a graph.
- `OracleEdge.setProperty`: Assigns a key-value property to an edge.

The following code fragment creates two vertices (`v1` and `v2`) and one edge (`e1`).

```
// Add vertices v1 and v2
Vertex v1 = opg.addVertex(11);
v1.setProperty("name", "Alice");
v1.setProperty("age", 31);

Vertex v2 = opg.addVertex(21);
v2.setProperty("name", "Bob");
v2.setProperty("age", 27);

// Add edge e1
Edge e1 = opg.addEdge(11, v1, v2, "knows");
e1.setProperty("type", "friends");
```

5.4.11 Deleting Vertices and Edges

You can remove vertex and edge instances individually, or all of them simultaneously. Use these methods:

- `OraclePropertyGraph.removeEdge`: Removes the specified edge from the graph.
- `OraclePropertyGraph.removeVertex`: Removes the specified vertex from the graph.
- `OraclePropertyGraph.clearRepository`: Removes all vertices and edges from the property graph instance.

The following code fragment removes edge `e1` and vertex `v1` from the graph instance. The adjacent edges will also be deleted from the graph when removing a vertex. This is because every edge must have an beginning and ending vertex. After removing the beginning or ending vertex, the edge is no longer a valid edge.

```
// Remove edge e1
opg.removeEdge(e1);

// Remove vertex v1
opg.removeVertex(v1);
```

The `OraclePropertyGraph.clearRepository` method can be used to remove all contents from an `OraclePropertyGraph` instance. However, use it with care because this action cannot be reversed.

5.4.12 Reading a Graph from a Database into an Embedded In-Memory Analyst

You can read a graph from Oracle Database into an in-memory analyst that is embedded in the same client Java application (a single JVM). For the following example, a correct `java.io.tmpdir` setting is required.

```
int dop = 8; // need customization
Map<PgxCfg.Field, Object> confPgxCfg = new HashMap<PgxCfg.Field,
Object>();
confPgxCfg.put(PgxCfg.Field.ENABLE_GM_COMPILER, false);
confPgxCfg.put(PgxCfg.Field.NUM_WORKERS_IO, dop); //
confPgxCfg.put(PgxCfg.Field.NUM_WORKERS_ANALYSIS, dop); // <= # of physical
cores
confPgxCfg.put(PgxCfg.Field.NUM_WORKERS_FAST_TRACK_ANALYSIS, 2);
confPgxCfg.put(PgxCfg.Field.SESSION_TASK_TIMEOUT_SECS, 0); // no timeout set
confPgxCfg.put(PgxCfg.Field.SESSION_IDLE_TIMEOUT_SECS, 0); // no timeout set
```

```
PgRdbmsGraphConfig cfg =
GraphConfigBuilder.forPropertyGraphRdbms().setJdbcUrl("jdbc:oracle:thin
:@<your_db_host>:<db_port>:<db_sid>")
    .setUsername("<username>")
    .setPassword("<password>")
    .setName("<graph_name>")
    .setMaxNumConnections(8)
    .setLoadEdgeLabel(false)
    .build();
OraclePropertyGraph opg = OraclePropertyGraph.getInstance(cfg);
ServerInstance localInstance = Pgx.getInstance();
localInstance.startEngine(confPgx);
PgxSession session = localInstance.createSession("session-id-1"); //
Put your session description here.

Analyst analyst = session.createAnalyst();

// The following call will trigger a read of graph data from the
database
PgxGraph pgxGraph = session.readGraphWithProperties(opg.getConfig());

long triangles = analyst.countTriangles(pgxGraph, false);
System.out.println("triangles " + triangles);

// Remove edge e1
opg.removeEdge(e1);

// Remove vertex v1
opg.removeVertex(v1);
```

5.4.13 Specifying Labels for Vertices

The database and data access layer do not provide labels for vertices; however, you can treat the value of a designated vertex property as one or more labels. Such a transformation is relevant only to the in-memory analyst.

In the following example, a property "country" is specified in a call to `setUseVertexPropertyValueAsLabel()`, and the comma delimiter "," is specified in a call to `setPropertyValueDelimiter()`. These two together imply that values of the country vertex property will be treated as vertex labels separated by a comma. For example, if vertex X has a string value "US" for its country property, then its vertex label will be US; and if vertex Y has a string value "UK,CN", then it will have two labels: UK and CN.

```
GraphConfigBuilder.forPropertyGraph...
    .setName("<your_graph_name>")
    ...
    .setUseVertexPropertyValueAsLabel("country")
    .setPropertyValueDelimiter(",")
    .setLoadVertexLabels(true)
    .build();
```

Related Topics

- [What Are Property Graphs?](#)

5.4.14 Building an In-Memory Graph

In addition to [Store the Database Password in a Keystore](#), you can create an in-memory graph programmatically. This can simplify development when the size of graph is small or when the content of the graph is highly dynamic. The key Java class is `GraphBuilder`, which can accumulate a set of vertices and edges added with the `addVertex` and `addEdge` APIs. After all changes are made, an in-memory graph instance (`PgxGraph`) can be created by the `GraphBuilder`.

The following Java code snippet illustrates a graph construction flow. Note that there are no explicit calls to `addVertex`, because any vertex that does not already exist will be added dynamically as its adjacent edges are created.

```
import oracle.pgx.api.*;

PgxSession session = Pgx.createSession("example");
GraphBuilder<Integer> builder = session.createGraphBuilder();

builder.addEdge(1, 2);
builder.addEdge(2, 3);
builder.addEdge(2, 4);
builder.addEdge(3, 4);
builder.addEdge(4, 2);

PgxGraph graph = builder.build();
```

To construct a graph with vertex properties, you can use `setProperty` against the vertex objects created.

```
PgxSession session = Pgx.createSession("example");
GraphBuilder<Integer> builder = session.createGraphBuilder();

builder.addVertex(1).setProperty("double-prop", 0.1);
builder.addVertex(2).setProperty("double-prop", 2.0);
builder.addVertex(3).setProperty("double-prop", 0.3);
builder.addVertex(4).setProperty("double-prop", 4.56789);

builder.addEdge(1, 2);
builder.addEdge(2, 3);
builder.addEdge(2, 4);
builder.addEdge(3, 4);
builder.addEdge(4, 2);

PgxGraph graph = builder.build();
```

To use long integers as vertex and edge identifiers, specify `IdType.LONG` when getting a new instance of `GraphBuilder`. For example:

```
import oracle.pgx.common.types.IdType;
GraphBuilder<Long> builder = session.createGraphBuilder(IdType.LONG);
```

During edge construction, you can directly use vertex objects that were previously created in a call to `addEdge`.

```
v1 = builder.addVertex(1).setProperty("double-prop", 0.5)
v2 = builder.addVertex(2).setProperty("double-prop", 2.0)

builder.addEdge(v1, v2)
```

As with vertices, edges can have properties. The following example sets the edge label by using `setLabel`:

```
builder.addEdge(v1, v2).setProperty("edge-prop",
"edge_prop_1_2").setLabel("label")
```

5.4.15 Dropping a Property Graph

To drop a property graph from the database, use the `OraclePropertyGraphUtils.dropPropertyGraph` method. This method has two parameters, the connection information and the graph name. For example:

```
// Drop the graph
Oracle oracle = new Oracle(jdbcUrl, username, password);
OraclePropertyGraphUtils.dropPropertyGraph(oracle, graphName);
```

You can also drop a property graph using the PL/SQL API. For example:

```
EXECUTE opg_apis.drop_pg('my_graph_name');
```

5.4.16 Executing PGQL Queries

You can execute PGQL queries directly against Oracle Database with the `PgqlStatement` and `PgqlPreparedStatement` interfaces. See [Executing PGQL Queries Against Property Graph Schema Tables](#) for details.

5.5 Managing Text Indexing for Property Graph Data

Indexes in Oracle Spatial and Graph property graph support allow fast retrieval of elements by a particular key/value or key/text pair. These indexes are created based on an element type (vertices or edges), a set of keys (and values), and an index type.

Oracle Spatial and Graph supports the use of the Oracle Text indexing technology, which is a feature of Oracle Database.

Two types of indexing structures are supported.

- Automatic text indexes provide automatic indexing of vertices or edges by a set of property keys. Their main purpose is to enhance query performance on vertices and edges based on particular key/value pairs.
- Manual text indexes enable you to define multiple indexes over a designated set of vertices and edges of a property graph. You must specify what graph elements go into the index.

Oracle Spatial and Graph provides APIs to create manual and automatic text indexes over property graphs stored in Oracle Database. Indexes are managed using Oracle Text, a proprietary search and analysis engine. The rest of this section focuses on how to create text indexes using the property graph capabilities of the Data Access Layer.

- [Configuring a Text Index for Property Graph Data](#)
- [Using Automatic Indexes for Property Graph Data](#)
- [Using Manual Indexes for Property Graph Data](#)
- [Executing Search Queries Over a Property Graph's Text Indexes](#)
- [Handling Data Types](#)
- [Updating Configuration Settings on Text Indexes for Property Graph Data](#)
Oracle's property graph support manages manual and automatic text indexes through integration with Oracle Text.
- [Using Parallel Query on Text Indexes for Property Graph Data](#)

5.5.1 Configuring a Text Index for Property Graph Data

The configuration of a text index is defined using an `OracleIndexParameters` object. This object includes information about the index such as search engine, location, number of directories (or shards), and degree of parallelism.

By default, text indexes are configured based on the `OracleIndexParameters` associated with the property graph using the method `opg.setDefaultIndexParameters(indexParams)`. The initial creation of the automatic index delimits the configuration and text search engine for future indexed keys.

Indexes can also be created by specifying a different set of parameters. The following code fragment creates an automatic text index over an existing property graph using a Lucene engine with a physical directory.

```
// Create an OracleIndexParameters object to get Index configuration (search engine,
etc).
OracleIndexParameters indexParams = OracleIndexParameters.buildFS(args)

// Create auto indexing on above properties for all vertices
opg.createKeyIndex("name", Vertex.class, indexParams.getParameters());
```

Any index configuration operations cause updates to be made to the IT\$ table, which is explained in [Property Graph Tables \(Detailed Information\)](#).

- [Configuring Text Indexes Using Oracle Text](#)

5.5.1.1 Configuring Text Indexes Using Oracle Text

Oracle Spatial and Graph supports automatic text indexes using Oracle Text. Oracle Text uses standard SQL to index, search, and analyze text values stored in the V column of the

vertices (or edges) table. Because Oracle Text indexes all the existing K/V pairs of the vertices (or edges) in the property graph, this option can be used **only** with automatic text indexes and must use a wildcard ("*") indexed key parameter during the index creation.

Because the property graph feature uses an NVARCHAR typed column for a better support of Unicode, it is highly recommended that UTF8 (AL32UTF8) be used as the database character set.

To create an Oracle Text index on the vertices table (or edges table), the ALTER SESSION privilege is required. The following example grants the privilege.

```
SQL> grant alter session to <YOUR_USER_SCHEMA_HERE>;
```

If customization is required, grant EXECUTE on CTX_DDL, as in the following example.

```
SQL> grant execute on ctx_ddl to <YOUR_USER_SCHEMA_HERE>;
```

A text index using Oracle Text uses an `OracleTextIndexParameters` object. The configuration parameters for indexes using a Oracle Text include:

- **Preference owner:** the owner of the preference.
- **Data store:** the datastore preference specifying how the text values are stored. A datastore preference can be created using `ctx_ddl.create_preference` API as follows:

```
SQL> -- The following requires access privilege to CTX_DDL
SQL> exec ctx_ddl.create_preference('SCOTT.OPG_DATASTORE',
'DIRECT_DATASTORE');
```

If the value is set to NULL, then the index will be created with `CTXSYS.DEFAULT_DATASORE`. This preference uses a `DIRECT_DATASTORE` type.

- **Filter:** the filter preference determining how text is filtered for indexing. A filter preference can be created using `ctx_ddl.create_preference`, as follows:

```
SQL> -- The following requires access privilege to CTX_DDL
SQL> exec ctx_ddl.create_preference('SCOTT.OPG_FILTER',
'AUTO_FILTER');
```

If the value is set to NULL, then the index will be created with `CTXSYS.NULL_FILTER`. This preference uses a `NULL_FILTER` type.

- **Storage:** the storage preference specifying table space and creation parameters for tables associated with a Text index. A storage preference can be created using `ctx_ddl.create_preference`, as follows:

```
SQL> -- The following requires access privilege to CTX_DDL
SQL> exec ctx_ddl.create_preference('SCOTT.OPG_STORAGE',
'BASIC_STORAGE');
```


If the value is set to NULL, then the index will be created with CTXSYS.DEFAULT_STORAGE. This preference uses a BASIC_STORAGE type.

- **Word list:** the word list preference specifying the enabled query options. These query options may include stemming, fuzzy matching, substring, and prefix indexing. A data store preference can be created using `ctx_ddl.create_preference`, as follows:

```
SQL> -- The following example enables stemming and fuzzy matching for
English.
SQL> exec ctx_ddl.create_preference('SCOTT.OPG_WORDLIST',
'BASIC_WORDLIST');
```

If the value is set to NULL, then the index will be created with CTXSYS.DEFAULT_WORDLIST. This preference uses the language stemmer for your database language.

- **Stop list:** the stop list preference specifying the list of words that are not meant to be indexed. A stop list preference can be created using `ctx_ddl.create_stoplist`.

If the value is set to NULL, then the index will be created with CTXSYS.DEFAULT_STOPLIST. This preference uses the stoplist of your database language.

- **Lexer:** the lexer preference specifying the language of the text to be indexed. A lexer preference can be created using `ctx_ddl.create_preference`, as follows:

```
SQL> -- The following requires access privilege to CTX_DDL
SQL> exec ctx_ddl.create_preference('SCOTT.OPG_AUTO_LEXER', 'AUTO_LEXER');
```

If the value is set to NULL, then the index will be created with CTXSYS.DEFAULT_LEXER. This preference uses a BASIC_LEXER type with additional options based on the language used at installation time.

The following code fragment creates the configuration for a text index using Oracle Text with default options and OPG_AUTO_LEXER.

```
String prefOwner = "scott";
String dataStore = (String) null;
String filter = (String) null;
String storage = (String) null;
String wordlist = (String) null;
String stoplist = (String) null;
String lexer = "OPG_AUTO_LEXER";
String options = (String) null;

OracleIndexParameters params
    = OracleTextIndexParameters.buildOracleText(prefOwner,
                                                dataStore,
                                                filter,
                                                storage,
                                                wordlist,
                                                stoplist,
                                                lexer,
                                                dop,
                                                options);
```

5.5.2 Using Automatic Indexes for Property Graph Data

An automatic text index provides automatic indexing of vertices or edges by a set of property keys. Its main purpose is to increase the speed of lookups over vertices and edges based on particular key/value pair. If an automatic index for the given key is enabled, then key/value pair lookups will be performed as a text search against the index instead of as a database lookup.

When specifying an automatic index over a property graph, use the following methods to create, remove, and manipulate an automatic index:

- `OraclePropertyGraph.createKeyIndex(String key, Class elementClass, Parameter[] parameters)`: Creates an automatic index for all elements of type `elementClass` by the given property key. The index is configured based on the specified parameters.
- `OraclePropertyGraph.createKeyIndex(String[] keys, Class elementClass, Parameter[] parameters)`: Creates an automatic index for all elements of type `elementClass` by using a set of property keys. The index is configured based on the specified parameters.
- `OraclePropertyGraph.dropKeyIndex(String key, Class elementClass)`: Drops the automatic index for all elements of type `elementClass` for the given property key.
- `OraclePropertyGraph.dropKeyIndex(String[] keys, Class elementClass)`: Drops the automatic index for all elements of type `elementClass` for the given set of property keys.
- `OraclePropertyGraph.getAutoIndex(Class elementClass)`: Gets an index instance of the automatic index for type `elementClass`.
- `OraclePropertyGraph.getIndexedKeys(Class elementClass)`: Gets the set of indexed keys currently used in an automatic index for all elements of type `elementClass`.

By default, indexes are configured based on the `OracleIndexParameters` associated with the property graph using the method `opg.setDefaultIndexParameters(indexParams)`.

Indexes can also be created by specifying a different set of parameters. This is shown in the following code snippet.

```
// Create an OracleIndexParameters object to get Index configuration (search engine, etc).
OracleIndexParameters indexParams = OracleIndexParameters.buildFS(args)

// Create auto indexing on above properties for all vertices
opg.createKeyIndex("name", Vertex.class, indexParams.getParameters());
```

The code fragment in the next example executes a query over all vertices to find all matching vertices with the key/value pair `name:Robert Smith`. This operation will execute a lookup into the text index.

Additionally, wildcard searches are supported by specifying the parameter `useWildCards` in the `getVertices` API call. Wildcard search is only supported when automatic indexes are enabled for the specified property key.

```
// Find all vertices with name Robert Smith.
Iterator<Vertices> vertices = opg.getVertices("name", "Robert Smith").iterator();
System.out.println("----- Vertices with name Robert Smith -----");
countV = 0;
while (vertices.hasNext()) {
    System.out.println(vertices.next());
    countV++;
}
System.out.println("Vertices found: " + countV);

// Find all vertices with name including keyword "Smith"
// Wildcard searching is supported.
boolean useWildcard = true;
Iterator<Vertices> vertices = opg.getVertices("name", "*Smith*").iterator();
System.out.println("----- Vertices with name *Smith* -----");
countV = 0;
while (vertices.hasNext()) {
    System.out.println(vertices.next());
    countV++;
}
System.out.println("Vertices found: " + countV);
```

The preceding code example produces output like the following:

```
----- Vertices with name Robert Smith-----
Vertex ID 1 {name:str:Robert Smith, role:str:political authority, occupation:str:CEO
of Example Corporation, country:str:United States, political party:str:Bipartisan,
religion:str:Unknown}
Vertices found: 1

----- Vertices with name *Smith* -----
Vertex ID 1 {name:str:Robert Smith, role:str:political authority, occupation:str:CEO
of Example Corporation, country:str:United States, political party:str:Bipartisan,
religion:str:Unknown}
Vertices found: 1
```

5.5.3 Using Manual Indexes for Property Graph Data

Manual indexes support the definition of multiple indexes over the vertices and edges of a property graph. A manual index requires that you manually put, get, and remove elements from the index.

When describing a manual index over a property graph, use the following methods to add, remove, and manipulate a manual index:

- `OraclePropertyGraph.createIndex(String name, Class elementClass, Parameter[] parameters)`: Creates a manual index with the specified name for all elements of type `elementClass`.
- `OraclePropertyGraph.dropIndex(String name)`: Drops the given manual index.
- `OraclePropertyGraph.getIndex(String name, Class elementClass)`: Gets an index instance of the given manual index for type `elementClass`.
- `OraclePropertyGraph.getIndices()`: Gets an array of index instances for all manual indexes created in the property graph.

5.5.4 Executing Search Queries Over a Property Graph's Text Indexes

Oracle Spatial and Graph provides a set of utilities to execute text search queries over automatic and manual text indexes. These utilities vary from querying based on a particular key/value pair, to executing a text search over a single or multiple keys (with extended query options such as wildcards, fuzzy searches, and range queries).

- [Executing Search Queries Over a Text Index Using Oracle Text](#)

5.5.4.1 Executing Search Queries Over a Text Index Using Oracle Text

Text search queries on Oracle Text are translated into SELECT SQL queries with a "contains" clause including a score range and ordering, and score ID. Oracle's property graph includes an utility called `OracleTextQueryObject`, which lets you execute text search queries over an Oracle Text index.

The following code fragment creates an automatic index using Oracle Text, and executes a query over the text index by specifying a particular key/value pair.

```
String prefOwner = "scott";
String datastore = (String) null;
String filter = (String) null;
String storage = (String) null;
String wordlist = (String) null;
String stoplist = (String) null;
String lexer = "OPG_AUTO_LEXER";
String options = (String) null;

OracleIndexParameters params
    =
OracleTextIndexParameters.buildOracleText(prefOwner,
                                           datastore,
                                           filter,
                                           storage,
                                           wordlist,
                                           stoplist,
                                           lexer,
                                           dop,
                                           options);

opg.setDefaultIndexParameters(indexParams);

// Create auto indexing on all existing properties, use wildcard for all
opg.createKeyIndex("*", Vertex.class);

// Get the auto index object
OracleIndex<Vertex> index = ((OracleIndex<Vertex>)
opg.getAutoIndex(Vertex.class));

// Create the text query object for Oracle Text
OracleTextQueryObject otqo
    = OracleTextQueryObject.getInstance("Smith" /* query body */,
                                       1 /* score */,
                                       ScoreRange.POSITIVE /* Score
range */,
                                       Direction.ASC /* order by
direction*/);
```

```

Iterator<Vertex> vertices = index.get("name", otqo).iterator();
System.out.println("----- Vertices with query: " + otqo.toString() + " -----");
countV = 0;
while (vertices.hasNext()) {
    System.out.println(vertices.next());
    countV++;
}
System.out.println("Vertices found: "+ countV);

```

You can filter the data type of the matching key/value pairs by specifying the data type class to execute the query against. The following code fragment executes a query over the text index to retrieve all properties with a String value including the word *Smith*.

```

// Create the text query object for Oracle Text
OracleTextQueryObject otqo
    = OracleTextQueryObject.getInstance("Smith" /* query body */,
                                        1 /* score */,
                                        ScoreRange.POSITIVE
                                        /* Score range */,
                                        Direction.ASC
                                        /* order by direction*/,
                                        "name",
                                        String.class);

Iterator<Vertex> vertices = index.get("name", otqo).iterator();
System.out.println("----- Vertices with query: " + otqo.toString() + " -----");
countV = 0;
while (vertices.hasNext()) {
    System.out.println(vertices.next());
    countV++;
}
System.out.println("Vertices found: "+ countV);

```

5.5.5 Handling Data Types

Oracle's property graph support indexes and stores an element's Key/Value pairs based on the value data type. The main purpose of handling data types is to provide extensive query support like numeric and date range queries.

By default, searches over a specific key/value pair are matched up to a query expression based on the value's data type. For example, to find vertices with the key/value pair `age:30`, a query is executed over all age fields with a data type integer. If the value is a query expression, you can also specify the data type class of the value to find by calling the API `get(String key, Object value, Class dtClass, Boolean useWildcards)`. If no data type is specified, the query expression will be matched to all possible data types.

When dealing with Boolean operators, each subsequent key/value pair must append the data type's prefix/suffix so the query can find proper matches.

- [Handling Data Types on Oracle Text](#)

5.5.5.1 Handling Data Types on Oracle Text

Text indexes using Oracle Text are created over the K and V text columns of the property graph tables. In order to provide text indexing capabilities on all available data types, Oracle populates the V column with a string representation of numeric, spatial, and date time key/value pairs.

To specify the date time and numeric formats used when populating the V column, you can use the methods `setNumberToCharSqlFormatString` and `setTimeToCharSqlFormatString`. The following code snippet shows how to set the date time and numeric formats in a property graph instance.

```
OraclePropertyGraph opg = OraclePropertyGraph.getInstance(args,
                                                         szGraphName);
opg.setNumberToCharSqlFormatString("TM9");
opg.setTimeToCharSqlFormatString("SYYYY-MM-DD\"T\"HH24:MI:SS.FF9TZH:TZM");
```

When executing a text search query over a numeric or date time value, you should use a text expression using the format associated to the property graph.

`OraclePropertyGraph` includes a utility API `opg.parseValueToCharSQLFormatString` that lets you parse a numeric or date time object into format used in the V column storage. The following code snippet calls this function with a date value and creates a text query object out of the retrieved text.

```
Date d = new java.util.Date(1001);
String szDate = opg.parseValueToCharSQLFormatString(d);

// Create the text query object for Oracle Text
OracleTextQueryObject otqo
    = OracleTextQueryObject.getInstance(szDate /* query body */,
                                       1 /* score */,
                                       ScoreRange.POSITIVE /* Score
range */,
                                       Direction.ASC /* order by
direction);
```

5.5.6 Updating Configuration Settings on Text Indexes for Property Graph Data

Oracle's property graph support manages manual and automatic text indexes through integration with Oracle Text.

At creation time, you must create an `OracleIndexParameters` object specifying the search engine and other configuration settings to be used by the text index. After a text index for property graph is created, these configuration settings cannot be changed.

For automatic indexes, all vertex index keys are managed by a single text index, and all edge index keys are managed by a different text index using the configuration specified when the first vertex or edge key is indexed.

If you need to change the configuration settings, you must first disable the current index and create it again using a new `OracleIndexParameters` object.

5.5.7 Using Parallel Query on Text Indexes for Property Graph Data

Text indexes in Oracle Spatial and Graph allow executing text queries over millions of vertices and edges by a particular key/value or key/text pair using parallel query execution.

Parallel text query will produce an array where each element holds all the vertices (or edges) with an attribute matching the given K/V pair from a shard. The subset of shards queried will be delimited by the given start sub-directory ID and the size of the connections array provided. This way, the subset will consider shards in the range of

[start, start - 1 + size of connections array]. Note that an integer ID (in the range of [0, N - 1]) is assigned to all the shards in index with N shards.

- [Parallel Text Search Using Oracle Text](#)

5.5.7.1 Parallel Text Search Using Oracle Text

You can use parallel text query using Oracle Text by calling the method `getPartitioned` in `OracleTextAutoIndex`, specifying an array of connections to Oracle Text (Connection objects), the key/value pair to search, and the starting partition ID.

The following code fragment generates an automatic text index using Oracle Text and executes a parallel text query. The number of calls to the `getPartitioned` method in the `OracleTextAutoIndex` class is controlled by the total number of partitions in the VT\$ (or GE\$ tables) and the number of connections used.

```
OraclePropertyGraph opg = OraclePropertyGraph.getInstance(...);
String prefOwner = "scott";
String datastore = (String) null;
String filter = (String) null;
String storage = (String) null;
String wordlist = (String) null;
String stoplist = (String) null;
String lexer = "OPG_AUTO_LEXER";
String options = (String) null;

OracleIndexParameters params
    = OracleTextIndexParameters.buildOracleText(prefOwner,
                                                datastore,
                                                filter,
                                                storage,
                                                wordlist,
                                                stoplist,
                                                lexer,
                                                dop,
                                                options);

opg.setDefaultIndexParameters(indexParams);

// Create auto indexing on all existing properties, use wildcard for all
opg.createKeyIndex("*", Vertex.class);

// Create the text query object for Oracle Text
OracleTextQueryObject otqo
    = OracleTextQueryObject.getInstance("Smith" /* query body */,
                                       1 /* score */,
                                       ScoreRange.POSITIVE /* Score range
*/,
                                       Direction.ASC /* order by
direction*/);

// Get the Connection object
Connection[] conns = new Connection[dop];
for (int idx = 0; idx < conns.length; idx++) {
    conns[idx] = opg.getOracle().clone().getConnection();
}

// Get the auto index object
```

```

OracleIndex<Vertex> index = ((OracleIndex<Vertex>)
opg.getAutoIndex(Vertex.class);

// Iterate to cover all the partitions in the index
long lCount = 0;
for (int split = 0; split < index.getTotalShards();
split += conns.length) {
    // Gets elements from split to split + conns.length
Iterable<Vertex>[] iterAr = index.getPartitioned(conns /* connections */,
"name"/* key */,
otqo,
true /* wildcards */,
split /* start split ID */);

lCount = countFromIterables(iterAr); /* Consume iterables in parallel */
}

// Close the connections
for (int idx = 0; idx < conns.length; idx++) {
conns[idx].dispose();
}

// Count results
System.out.println("Vertices found using parallel query: " + lCount);

```

5.6 Access Control for Property Graph Data (Graph-Level and OLS)

Oracle Graph supports two access control and security models: graph level access control, and fine-grained security through integration with Oracle Label Security (OLS).

- Graph-level access control relies on grant/revoke to allow/disallow users other than the owner to access a property graph.
- OLS for property graph data allows sensitivity labels to be associated with individual vertex or edge stored in a property graph.

The default control of access to property graph data stored in an Oracle Database is at the graph level: the owner of a graph can grant read, insert, delete, update and select privileges on the graph to other users.

However, for applications with stringent security requirements, you can enforce a fine-grained access control mechanism by using the Oracle Label Security option of Oracle Database. With OLS, for each query, access to specific elements (vertices or edges) is granted by comparing their labels with the user's labels. (For information about using OLS, see *Oracle Label Security Administrator's Guide* .)

With Oracle Label Security enabled, elements (vertices or edges) may not be inserted in the graph if the same elements exist in the database with a stronger sensitivity label. For example, assume that you have a vertex with a very sensitive label, such as: (Vertex ID 1 {name:str:v1} "SENSITIVE"). This actually prevents a low-privileged (PUBLIC) user from updating the vertex: (Vertex ID 1 {name:str:v1} "PUBLIC"). On the other hand, if a high-privileged user overwrites a vertex or an edge that had been created with a low-level security label, the newer label with higher security will be assigned to the vertex or edge, and the low-privileged user will not be able to see it anymore.

- [Applying Oracle Label Security \(OLS\) on Property Graph Data](#)
This topic presents an example illustrating how to apply OLS to property graph data.

5.6.1 Applying Oracle Label Security (OLS) on Property Graph Data

This topic presents an example illustrating how to apply OLS to property graph data.

Because the property graph is stored in regular relational tables, this example is no different from applying OLS on a regular relational table. The following shows how to configure and enable OLS, create a security policy with security labels, and apply it to a property graph. The code examples are very simplified, and do not necessarily reflect recommended practices regarding user names and passwords.

1. As SYSDBA, create database users named userP, userP2, userS, userTS, userTS2 and pgAdmin.

```
CONNECT / as sysdba;
```

```
CREATE USER userP IDENTIFIED BY userPpass;  
GRANT connect, resource, create table, create view, create any index TO  
userP;  
GRANT unlimited TABLESPACE to userP;
```

```
CREATE USER userP2 IDENTIFIED BY userP2pass;  
GRANT connect, resource, create table, create view, create any index TO  
userP2;  
GRANT unlimited TABLESPACE to userP2;
```

```
CREATE USER userS IDENTIFIED BY userSpass;  
GRANT connect, resource, create table, create view, create any index TO  
userS;  
GRANT unlimited TABLESPACE to userS;
```

```
CREATE USER userTS IDENTIFIED BY userTSpass;  
GRANT connect, resource, create table, create view, create any index TO  
userTS;  
GRANT unlimited TABLESPACE to userTS;
```

```
CREATE USER userTS2 IDENTIFIED BY userTS2pass;  
GRANT connect, resource, create table, create view, create any index TO  
userTS2;  
GRANT unlimited TABLESPACE to userTS2;
```

```
CREATE USER pgAdmin IDENTIFIED BY pgAdminpass;  
GRANT connect, resource, create table, create view, create any index TO  
pgAdmin;  
GRANT unlimited TABLESPACE to pgAdmin;
```

2. As SYSDBA, configure and enable Oracle Label Security.

```
ALTER USER lbacsys IDENTIFIED BY lbacsys ACCOUNT UNLOCK;  
EXEC LBACSYS.CONFIGURE_OLS;  
EXEC LBACSYS.OLS_ENFORCEMENT.ENABLE_OLS;
```

3. As SYSTEM, grant privileges to sec_admin and hr_sec.

```
CONNECT system/<system-password>
GRANT connect, create any index to sec_admin IDENTIFIED BY password;
GRANT connect, create user, drop user, create role, drop any role
TO hr_sec IDENTIFIED BY password;
```

4. As LBACSYS, create the security policy.

```
CONNECT lbacsys/<lbacsys-password>

BEGIN
SA_SYSDBA.CREATE_POLICY (
  policy_name => 'DEFENSE',
  column_name => 'SL',
  default_options => 'READ_CONTROL, LABEL_DEFAULT, HIDE');
END;
/
```

5. As LBACSYS , grant DEFENSE_DBA and execute to sec_admin and hr_sec users.

```
GRANT DEFENSE_DBA to sec_admin;
GRANT DEFENSE_DBA to hr_sec;

GRANT execute on SA_COMPONENTS to sec_admin;
GRANT execute on SA_USER_ADMIN to hr_sec;
```

6. As SEC_ADMIN, create three security levels (For simplicity, compartments and groups are omitted here.)

```
CONNECT sec_admin/<sec_admin-password>;

BEGIN
SA_COMPONENTS.CREATE_LEVEL (
  policy_name => 'DEFENSE',
  level_num => 1000,
  short_name => 'PUB',
  long_name => 'PUBLIC');
END;
/
EXECUTE
SA_COMPONENTS.CREATE_LEVEL('DEFENSE',2000,'CONF','CONFIDENTIAL');
EXECUTE
SA_COMPONENTS.CREATE_LEVEL('DEFENSE',3000,'SENS','SENSITIVE');
```

7. Create three labels.

```
EXECUTE SA_LABEL_ADMIN.CREATE_LABEL('DEFENSE',1000,'PUB');
EXECUTE SA_LABEL_ADMIN.CREATE_LABEL('DEFENSE',2000,'CONF');
EXECUTE SA_LABEL_ADMIN.CREATE_LABEL('DEFENSE',3000,'SENS');
```

8. As HR_SEC, assign labels and privileges.

```
CONNECT hr_sec/<hr_sec-password>;

BEGIN
SA_USER_ADMIN.SET_USER_LABELS (
  policy_name => 'DEFENSE',
  user_name => 'UT',
  max_read_label => 'SENS',
  max_write_label => 'SENS',
  min_write_label => 'CONF',
  def_label => 'SENS',
  row_label => 'SENS');
END;
/

EXECUTE SA_USER_ADMIN.SET_USER_LABELS('DEFENSE', 'userTS', 'SENS');
EXECUTE SA_USER_ADMIN.SET_USER_LABELS('DEFENSE', 'userTS2', 'SENS');
EXECUTE SA_USER_ADMIN.SET_USER_LABELS('DEFENSE', 'userS', 'CONF');
EXECUTE SA_USER_ADMIN.SET_USER_LABELS ('DEFENSE', userP, 'PUB', 'PUB',
  'PUB', 'PUB', 'PUB');
EXECUTE SA_USER_ADMIN.SET_USER_LABELS ('DEFENSE', 'userP2', 'PUB', 'PUB',
  'PUB', 'PUB', 'PUB');
EXECUTE SA_USER_ADMIN.SET_USER_PRIVS ('DEFENSE', 'pgAdmin', 'FULL');
```

9. As SEC_ADMIN, apply the security policies to the desired property graph. Assume a property graph with the name OLSEXAMPLE with userP as the graph owner. To apply OLS security, execute the following statements.

```
CONNECT sec_admin/<password>;

EXECUTE SA_POLICY_ADMIN.APPLY_TABLE_POLICY ('DEFENSE', 'userP',
  'OLSEXAMPLEVT$');
EXECUTE SA_POLICY_ADMIN.APPLY_TABLE_POLICY ('DEFENSE', 'userP',
  'OLSEXAMPLEGE$');
EXECUTE SA_POLICY_ADMIN.APPLY_TABLE_POLICY ('DEFENSE', 'userP',
  'OLSEXAMPLEGT$');
EXECUTE SA_POLICY_ADMIN.APPLY_TABLE_POLICY ('DEFENSE', 'userP',
  'OLSEXAMPLESS$');
```

Now Oracle Label Security has sensitivity labels to be associated with individual vertices or edges stored in the property graph.

The following example shows how to create a property graph with name OLSEXAMPLE, and an example flow to demonstrate the behavior when different users with different security labels create, read, and write graph elements.

```
// Create Oracle Property Graph
String graphName = "OLSEXAMPLE";
Oracle connPub = new Oracle("jdbc:oracle:thin:@host:port:SID", "userP",
  "userPpass");
OraclePropertyGraph graphPub = OraclePropertyGraph.getInstance(connPub,
  graphName, 48);

// Grant access to other users
```

```

graphPub.grantAccess("userP2", "RSIUD"); // Read, Select, Insert,
Update, Delete (RSIUD)
graphPub.grantAccess("userS", "RSIUD");
graphPub.grantAccess("userTS", "RSIUD");
graphPub.grantAccess("userTS2", "RSIUD");

// Load data
OraclePropertyGraphDataLoader opgdl =
OraclePropertyGraphDataLoader.getInstance();
String vfile = "../data/connections.opv";
String efile = "../data/connections.ope";
graphPub.clearRepository();
opgdl.loadData(graphPub, vfile, efile, 48, 1000, true, null);
System.out.println("Vertices with user userP and PUBLIC LABEL: " +
graphPub.countVertices()); // 78
System.out.println("Vertices with user userP and PUBLIC LABEL: " +
graphPub.countEdges()); // 164

// Second user with a higher level
Oracle connTS = new Oracle("jdbc:oracle:thin:@host:port:SID",
"userTS", "userTpassS");
OraclePropertyGraph graphTS = OraclePropertyGraph.getInstance(connTS,
"USERP", graphName, 8, 48, null, null);
System.out.println("Vertices with user userTS and SENSITIVE LABEL: " +
graphTS.countVertices()); // 78
System.out.println("Vertices with user userTS and SENSITIVE LABEL: " +
graphTS.countEdges()); // 164

// Add vertices and edges with the second user
long lMaxVertexID = graphTS.getMaxVertexID();
long lMaxEdgeID = graphTS.getMaxEdgeID();
long size = 10;
System.out.println("\nAdd " + size + " vertices and edges with user
userTS and SENSITIVE LABEL\n");
for (long idx = 1; idx <= size; idx++) {
    Vertex v = graphTS.addVertex(idx + lMaxVertexID);
    v.setProperty("name", "v_" + (idx + lMaxVertexID));
    Edge e = graphTS.addEdge(idx + lMaxEdgeID, v,
graphTS.getVertex(idx), "edge_" + (idx + lMaxEdgeID));
}
graphTS.commit();

// User userP with a lower level only sees the original vertices and
edges, user userTS can see more
System.out.println("Vertices with user userP and PUBLIC LABEL: " +
graphPub.countVertices()); // 78
System.out.println("Vertices with user userP and PUBLIC LABEL: " +
graphPub.countEdges()); // 164
System.out.println("Vertices with user userTS and SENSITIVE LABEL: " +
graphTS.countVertices()); // 88
System.out.println("Vertices with user userTS and SENSITIVE LABEL: " +
graphTS.countEdges()); // 174

// Third user with a higher level
Oracle connTS2 = new Oracle("jdbc:oracle:thin:@host:port:SID",

```

```

"userTS2", "userTS2pass");
OraclePropertyGraph graphTS2 = OraclePropertyGraph.getInstance(connTS2,
"USERP", graphName, 8, 48, null, null);
System.out.println("Vertices with user userTS2 and SENSITIVE LABEL: " +
graphTS2.countVertices()); // 88
System.out.println("Vertices with user userTS2 and SENSITIVE LABEL: " +
graphTS2.countEdges()); // 174

// Fourth user with a intermediate level
Oracle connS = new Oracle("jdbc:oracle:thin:@host:port:SID", "userS",
"userSpass");
OraclePropertyGraph graphS = OraclePropertyGraph.getInstance(connS, "USERP",
graphName, 8, 48, null, null);
System.out.println("Vertices with user userS and CONFIDENTIAL LABEL: " +
graphS.countVertices()); // 78
System.out.println("Vertices with user userS and CONFIDENTIAL LABEL: " +
graphS.countEdges()); // 164

// Modify vertices with the fourth user
System.out.println("\nModify " + size + " vertices with user userS and
CONFIDENTIAL LABEL\n");
for (long idx = 1; idx <= size; idx++) {
    Vertex v = graphS.getVertex(idx);
    v.setProperty("security_label", "CONFIDENTIAL");
}
graphS.commit();

// User userP with a lower level that userS cannot see the new vertices
// Users userS and userTS can see them
System.out.println("Vertices with user userP with property security_label: "
+ OraclePropertyGraphUtils.size(graphPub.getVertices("security_label",
"CONFIDENTIAL"))); // 0
System.out.println("Vertices with user userS with property security_label: "
+ OraclePropertyGraphUtils.size(graphS.getVertices("security_label",
"CONFIDENTIAL"))); // 10
System.out.println("Vertices with user userTS with property security_label:
" + OraclePropertyGraphUtils.size(graphTS.getVertices("security_label",
"CONFIDENTIAL"))); // 10
System.out.println("Vertices with user userP and PUBLIC LABEL: " +
graphPub.countVertices()); // 68
System.out.println("Vertices with user userTS and SENSITIVE LABEL: " +
graphTS.countVertices()); // 88

```

The preceding example should produce the following output.

```

Vertices with user userP and PUBLIC LABEL: 78
Vertices with user userP and PUBLIC LABEL: 164
Vertices with user userTS and SENSITIVE LABEL: 78
Vertices with user userTS and SENSITIVE LABEL: 164

```

Add 10 vertices and edges with user userTS and SENSITIVE LABEL

```

Vertices with user userP and PUBLIC LABEL: 78
Vertices with user userP and PUBLIC LABEL: 164
Vertices with user userTS and SENSITIVE LABEL: 88
Vertices with user userTS and SENSITIVE LABEL: 174

```

```

Vertices with user userTS2 and SENSITIVE LABEL: 88
Vertices with user userTS2 and SENSITIVE LABEL: 174
Vertices with user userS and CONFIDENTIAL LABEL: 78
Vertices with user userS and CONFIDENTIAL LABEL: 164

```

```

Modify 10 vertices with user userS and CONFIDENTIAL LABEL

```

```

Vertices with user userP with property security_label: 0
Vertices with user userS with property security_label: 10
Vertices with user userTS with property security_label: 10
Vertices with user userP and PUBLIC LABEL: 68
Vertices with user userTS and SENSITIVE LABEL: 88

```

5.7 SQL-Based Property Graph Query and Analytics

You can use SQL to query property graph data in Oracle Spatial and Graph.

For the property graph support in Oracle Spatial and Graph, all the vertices and edges data are persisted in relational form in Oracle Database. For detailed information about the Oracle Spatial and Graph property graph schema objects, see [Property Graph Schema Objects for Oracle Database](#).

This chapter provides examples of typical graph queries implemented using SQL. The audience includes DBAs as well as application developers who understand SQL syntax and property graph schema objects.

The benefits of querying directly property graph using SQL include:

- There is no need to bring data outside Oracle Database.
- You can leverage the industry-proven SQL engine provided by Oracle Database.
- You can easily join or integrate property graph data with other data types (relational, JSON, XML, and so on).
- You can take advantage of existing Oracle SQL tuning and database management tools and user interface.

The examples assume that there is a property graph named `connections` in the current schema. The SQL queries and example output are for illustration purpose only, and your output may be different depending on the data in your `connections` graph. In some examples, the output is reformatted for readability.

- [Simple Property Graph Queries](#)
The examples in this topic query vertices, edges, and properties of the graph.
- [Text Queries on Property Graphs](#)
If values of a property (vertex property or edge property) contain free text, then it might help performance to create an Oracle Text index on the V column.
- [Navigation and Graph Pattern Matching](#)
A key benefit of using a graph data model is that you can easily navigate across entities (people, movies, products, services, events, and so on) that are modeled as vertices, following links and relationships modeled as edges. In addition, graph matching templates can be defined to do such things as detect patterns, aggregate individuals, and analyze trends.
- [Navigation Options: CONNECT BY and Parallel Recursion](#)
The CONNECT BY clause and parallel recursion provide options for advanced navigation and querying.

- [Pivot](#)
The PIVOT clause lets you dynamically add columns to a table to create a new table.
- [SQL-Based Property Graph Analytics](#)
In addition to the analytical functions offered by the in-memory analyst, the property graph feature in Oracle Spatial and Graph supports several native, SQL-based property graph analytics.

5.7.1 Simple Property Graph Queries

The examples in this topic query vertices, edges, and properties of the graph.

Example 5-1 Find a Vertex with a Specified Vertex ID

This example find the vertex with vertex ID 1 in the `connections` graph.

```
SQL> select vid, k, v, vn, vt
       from connectionsVT$
       where vid=1;
```

The output might be as follows:

```
1 country      United States
1 name         Robert Smith
1 occupation   CEO of Example Corporation
...
```

Example 5-2 Find an Edge with a Specified Edge ID

This example find the edge with edge ID 100 in the `connections` graph.

```
SQL> select eid,svid,dvid,k,t,v,vn,vt
       from connectionsGE$
       where eid=1000;
```

The output might be as follows:

```
1000 1 2 weight 3 1 1
```

In the preceding output, the K of the edge property is "weight" and the type ID of the value is 3, indicating a float value.

Example 5-3 Perform Simple Counting

This example performs simple counting in the `connections` graph.

```
SQL> -- Get the total number of K/V pairs of all the vertices
SQL> select /*+ parallel */ count(1)
       from connectionsVT$;
```

```
299
```

```
SQL> -- Get the total number of K/V pairs of all the edges
SQL> select /*+ parallel(8) */ count(1)
       from connectionsGE$;
```

```
164
```

```
SQL> -- Get the total number of vertices
SQL> select /*+ parallel */ count(distinct vid)
       from connectionsVT$;
```

78

```
SQL> -- Get the total number of edges
SQL> select /*+ parallel */ count(distinct eid)
       from connectionsGE$;
```

164

Example 5-4 Get the Set of Property Keys Used

This example gets the set of property keys used for the vertices in the `connections` graph.

```
SQL> select /*+ parallel */ distinct k
       from connectionsVT$;
```

```
company
show
occupation
type
team
religion
criminal charge
music genre
genre
name
role
political party
country
```

13 rows selected.

```
SQL> -- get the set of property keys used for edges
SQL> select /*+ parallel */ distinct k
       from connectionsGE$;
```

```
weight
```

Example 5-5 Find Vertices with a Value

This example finds vertices with a value (of any property) that is of String type, and where the value contains two adjacent occurrences of a, e, i, o, or u, regardless of case in the `connections` graph.

```
SQL> select vid, t, k, v
       from connectionsVT$
       where t=1
          and regexp_like(v, '([aeiou])\1', 'i');
```



```

6          1 name Jordan Peele
6          1 show Key and Peele
54         1 name John Green
...

```

It is usually hard to leverage a B-Tree index for the preceding kind of query because it is difficult to know beforehand what kind of regular expression is going to be used. For the above query, you might get the following execution plan. Note that full table scan is chosen by the optimizer.

```

-----
| Id | Operation          | Name          | Rows | Bytes | Cost (%CPU)| Time
| Pstart| Pstop |   TQ   | IN-OUT| PQ Distrib |
-----
|  0 | SELECT STATEMENT   |               |    15 |   795 |    28  (0)| 00:00:01
|      |               |               |      |      |          |
|  1 | PX COORDINATOR     |               |      |      |          |
|      |               |               |      |      |          |
|  2 | PX SEND QC (RANDOM)| :TQ10000     |    15 |   795 |    28  (0)| 00:00:01
|      |      | Q1,00 | P->S | QC (RAND) |
|  3 | PX BLOCK ITERATOR |               |    15 |   795 |    28  (0)| 00:00:01
|      |      |      |      |           |
|*  4 |  TABLE ACCESS FULL| CONNECTIONSVT$ |    15 |   795 |    28  (0)| 00:00:01
|      |      |      |      |           |
-----

```

Predicate Information (identified by operation id):

```

4 - filter(INTERNAL_FUNCTION("V") AND REGEXP_LIKE ("V",U'([aeiou])\005C1','i') AND
"T"=1 AND INTERNAL_FUNCTION("K"))

```

Note

- Degree of Parallelism is 2 because of table property

If the Oracle Database In-Memory option is available and memory is sufficient, it can help performance to place the table (full table or a set of relevant columns) in memory. One way to achieve that is as follows:

```

SQL> alter table connectionsVT$ inmemory;
Table altered.

```

Now, entering the same SQL containing the regular expression shows a plan that performs a "TABLE ACCESS INMEMORY FULL".

```

-----
| Id | Operation          | Name          | Rows | Bytes | Cost (%CPU)| Time
| Time   | Pstart| Pstop |   TQ   | IN-OUT| PQ Distrib |
-----
|  0 | SELECT STATEMENT   |               |    15 |   795 |    28
(0)| 00:00:01 |      |      |      |      |
|  1 | PX COORDINATOR     |               |      |      |          |
|      |               |               |      |      |          |
|  2 | PX SEND QC (RANDOM)| :TQ10000     |    15 |   795 |    28
(0)| 00:00:01 |      | Q1,00 | P->S | QC (RAND) |
-----

```

```

| 3 | PX BLOCK ITERATOR | 15 | 795 |
28 (0) | 00:00:01 | 1 | 8 | Q1,00 | PCWC |
|* 4 | TABLE ACCESS INMEMORY FULL | CONNECTIONSVT$ | 15 | 795 |
28 (0) | 00:00:01 | 1 | 8 | Q1,00 | PCWP |
-----

```

```

-----
Predicate Information (identified by operation id):
-----

```

```

4 - filter(INTERNAL_FUNCTION("V") AND REGEXP_LIKE ("V",U'([aeiou])
\005C1','i') AND "T"=1 AND INTERNAL_FUNCTION("K"))

```

```

Note
-----

```

```

- Degree of Parallelism is 2 because of table property

```

5.7.2 Text Queries on Property Graphs

If values of a property (vertex property or edge property) contain free text, then it might help performance to create an Oracle Text index on the V column.

Oracle Text can process text that is directly stored in the database. The text can be short strings (such as names or addresses), or it can be full-length documents. These documents can be in a variety of textual format.

The text can also be in many different languages. Oracle Text can handle any space-separated languages (including character sets such as Greek or Cyrillic). In addition, Oracle Text is able to handle the Chinese, Japanese and Korean pictographic languages)

Because the property graph feature uses NVARCHAR typed column for better support of Unicode, it is **highly recommended** that UTF8 (AL32UTF8) be used as the database character set.

To create an Oracle Text index on the vertices table (or edges table), the ALTER SESSION privilege is required. For example:

```
SQL> grant alter session to <YOUR_USER_SCHEMA_HERE>;
```

If customization is required, also grant the EXECUTE privilege on CTX_DDL:

```
SQL> grant execute on ctx_ddl to <YOUR_USER_SCHEMA_HERE>;
```

The following shows some example statements for granting these privileges to SCOTT.

```

SQL> conn / as sysdba
Connected.
SQL> -- This is a PDB setup --
SQL> alter session set container=orcl;
Session altered.

SQL> grant execute on ctx_ddl to scott;
Grant succeeded.

SQL> grant alter session to scott;
Grant succeeded.

```

Example 5-6 Create a Text Index

This example creates an Oracle Text index on the vertices table (V column) of the connections graph in the SCOTT schema. Note that the Oracle Text index created here is for all property keys, not just one or a subset of property keys. In addition, if a new property is added to the graph and the property value is of String data type, then it will automatically be included in the same text index.

The example uses the OPG_AUTO_LEXER lexer owned by MDSYS.

```
SQL> execute opg_apis.create_vertices_text_idx('scott', 'connections',
pref_owner=>'MDSYS', lexer=>'OPG_AUTO_LEXER', dop=>2);
```

If customization is desired, you can use the `ctx_ddl.create_preference` API. For example:

```
SQL> -- The following requires access privilege to CTX_DDL
SQL> exec ctx_ddl.create_preference('SCOTT.OPG_AUTO_LEXER', 'AUTO_LEXER');
```

PL/SQL procedure successfully completed.

```
SQL> execute opg_apis.create_vertices_text_idx('scott', 'connections',
pref_owner=>'scott', lexer=>'OPG_AUTO_LEXER', dop=>2);
```

PL/SQL procedure successfully completed.

You can now use a rich set of functions provided by Oracle Text to perform queries against graph elements.

Note:

If you no longer need an Oracle Text index, you can use the `drop_vertices_text_idx` or `opg_apis.drop_edges_text_idx` API to drop it. The following statements drop the text indexes on the vertices and edges of a graph named `connections` owned by SCOTT:

```
SQL> exec opg_apis.drop_vertices_text_idx('scott', 'connections');
SQL> exec opg_apis.drop_edges_text_idx('scott', 'connections');
```

Example 5-7 Find a Vertex that Has a Property Value

The following example find a vertex that has a property value (of string type) containing the keyword "Smith".

```
SQL> select vid, k, t, v
       from connectionsVT$
       where t=1
          and contains(v, 'Smith', 1) > 0
       order by score(1) desc
       ;
```

The output and SQL execution plan from the preceding statement may appear as follows. Note that DOMAIN INDEX appears as an operation in the execution plan.

```

1 name      1 Robert Smith

Execution Plan
-----
Plan hash value: 1619508090

-----
| Id | Operation                               | Name                | Rows | Bytes |
Cost (%CPU) | Time      | Pstart| Pstop |
-----
|  0 | SELECT STATEMENT                       |                    |    1 |    56 |
|  5 |   (20) | 00:00:01 |      |      |
|  1 |    SORT ORDER BY                       |                    |    1 |    56 |
|  5 |   (20) | 00:00:01 |      |      |
|*  2 |    TABLE ACCESS BY GLOBAL INDEX ROWID| CONNECTIONSVT$     |    1 |    56 |
|  4 |       (0) | 00:00:01 | ROWID | ROWID |
|*  3 |      DOMAIN INDEX                       | CONNECTIONSXTV$    |      |      |
|  4 |       (0) | 00:00:01 |      |      |
-----

Predicate Information (identified by operation id):
-----

2 - filter("T"=1 AND INTERNAL_FUNCTION("K") AND INTERNAL_FUNCTION("V"))
3 - access("CTXSYS"."CONTAINS"("V",'Smith',1)>0)

```

Example 5-8 Fuzzy Match

The following example finds a vertex that has a property value (of string type) containing variants of "ameriian" (a deliberate misspelling for this example) Fuzzy match is used.

```

SQL> select vid, k, t, v
       from connectionsVT$
       where contains(v, 'fuzzy(ameriian,,weight)', 1) > 0
       order by score(1) desc;

```

The output and SQL execution plan from the preceding statement may appear as follows.

```

8 role      1 american business man
9 role      1 american business man
4 role      1 american economist
6 role      1 american comedian actor
7 role      1 american comedian actor
1 occupation 1 44th president of United States of America

6 rows selected.

Execution Plan
-----
Plan hash value: 1619508090

-----

```

```
-----
| Id | Operation                                | Name                    | Rows | Bytes | Cost
(%CPU)| Time          | Pstart| Pstop |      |      |
-----
|  0 | SELECT STATEMENT                        |                         |     1 |    56 |    5
(20)| 00:00:01 |      |      |      |      |
|  1 | SORT ORDER BY                          |                         |     1 |    56 |    5
(20)| 00:00:01 |      |      |      |      |
|*  2 | TABLE ACCESS BY GLOBAL INDEX ROWID| CONNECTIONSVT$         |     1 |    56 |
4   (0)| 00:00:01 | ROWID | ROWID |
|*  3 | DOMAIN INDEX                            | CONNECTIONSXTV$       |      |      |
4   (0)| 00:00:01 |      |      |      |      |
-----
```

Predicate Information (identified by operation id):

```
-----
2 - filter(INTERNAL_FUNCTION("K") AND INTERNAL_FUNCTION("V"))
```

Example 5-9 Query Relaxation

The following example is a sophisticated Oracle Text query that implements **query relaxation**, which enables you to execute the most restrictive version of a query first, progressively relaxing the query until the required number of matches is obtained. Using query relaxation with queries that contain multiple strings, you can provide guidance for determining the “best” matches, so that these appear earlier in the results than other potential matches.

This example searches for "american actor" with a query relaxation sequence.

```
SQL> select vid, k, t, v
       from connectionsVT$
       where CONTAINS (v,
'<query>
<textquery lang="ENGLISH" grammar="CONTEXT">
  <progression>
    <seq>{american} {actor}</seq>
    <seq>{american} NEAR {actor}</seq>
    <seq>{american} AND {actor}</seq>
    <seq>{american} ACCUM {actor}</seq>
  </progression>
</textquery>
<score datatype="INTEGER" algorithm="COUNT"/>
</query>') > 0;
```

The output and SQL execution plan from the preceding statement may appear as follows.

```
7 role      1 american comedian actor
6 role      1 american comedian actor
44 occupation 1 actor
8 role      1 american business man
53 occupation 1 actor film producer
52 occupation 1 actor
4 role      1 american economist
47 occupation 1 actor
9 role      1 american business man
```

9 rows selected.

Execution Plan

Plan hash value: 2158361449

Id	Operation	Name	Rows	Bytes	Cost
(%CPU)	Time				
0	SELECT STATEMENT		1		56
4	(0) 00:00:01				
* 1	TABLE ACCESS BY GLOBAL INDEX ROWID	CONNECTIONSVT\$	1		56
4	(0) 00:00:01 ROWID ROWID				
* 2	DOMAIN INDEX	CONNECTIONSXTV\$			
4	(0) 00:00:01				

Predicate Information (identified by operation id):

1 - filter(INTERNAL_FUNCTION("K") AND INTERNAL_FUNCTION("V"))
2 - access("CTXSYS"."CONTAINS"("V",'<query> <textquery lang="ENGLISH" grammar="CONTEXT">
<progression> <seq>{american} {actor}</seq> <seq>{american}
NEAR {actor}</seq>
<seq>{american} AND {actor}</seq> <seq>{american} ACCUM
{actor}</seq> </progression>
</textquery> <score datatype="INTEGER" algorithm="COUNT"/> </
query>')>0)

Example 5-10 Find an Edge

Just as with vertices, you can create an Oracle Text index on the V column of the edges table (GE\$) of a property graph. The following example uses the OPG_AUTO_LEXER lexer owned by MDSYS.

```
SQL> exec opg_apis.create_edges_text_idx('scott', 'connections',
pref_owner=>'mdsys', lexer=>'OPG_AUTO_LEXER', dop=>4);
```

If customization is required, use the `ctx_ddl.create_preference` API.

5.7.3 Navigation and Graph Pattern Matching

A key benefit of using a graph data model is that you can easily navigate across entities (people, movies, products, services, events, and so on) that are modeled as vertices, following links and relationships modeled as edges. In addition, graph matching templates can be defined to do such things as detect patterns, aggregate individuals, and analyze trends.

This topic provides graph navigation and pattern matching examples using the example property graph named connections. Most of the SQL statements are relatively simple, but they can be used as building blocks to implement requirements

that are more sophisticated. It is generally best to start from something simple, and progressively add complexity.

Example 5-11 Who Are a Person's Collaborators?

The following SQL statement finds all entities that a vertex with ID 1 collaborates with. For simplicity, it considers **only** outgoing relationships.

```
SQL> select dvid, el, k, vn, v
       from connectionsGE$
       where svid=1
          and el='collaborates';
```



Note:

To find the specific vertex ID of interest, you can perform a text query on the property graph using keywords or fuzzy matching. (For details and examples, see [Text Queries on Property Graphs](#).)

The preceding example's output and execution plan may be as follows.

```
2 collaborates weight 1 1
21 collaborates weight 1 1
22 collaborates weight 1 1
....
26 collaborates weight 1 1
```

10 rows selected.

Id	Operation	Pstart	Pstop	TQ	Name	Rows	Bytes
Cost (%CPU)	Time				IN-OUT PQ Distrib		
0	SELECT STATEMENT					10	460
2	(0) 00:00:01						
1	PX COORDINATOR						
2	PX SEND QC (RANDOM)				:TQ10000	10	460
2	(0) 00:00:01			Q1,00	P->S QC (RAND)		
3	PX PARTITION HASH ALL					10	460
2	(0) 00:00:01	1	8	Q1,00	PCWC		
* 4	TABLE ACCESS BY LOCAL INDEX ROWID BATCHED				CONNECTIONSGE\$	10	460
2	(0) 00:00:01	1	8	Q1,00	PCWP		
* 5	INDEX RANGE SCAN				CONNECTIONSXSE\$	20	
1	(0) 00:00:01	1	8	Q1,00	PCWP		

Predicate Information (identified by operation id):

```
4 - filter(INTERNAL_FUNCTION("EL") AND "EL"=U'collaborates' AND
```

```
INTERNAL_FUNCTION("K") AND INTERNAL_FUNCTION("V"))
5 - access("SVID"=1)
```

Example 5-12 Who Are a Person's Collaborators and What are Their Occupations?

The following SQL statement finds collaborators of the vertex with ID 1, and the occupation of each collaborator. A join with the vertices table (VT\$) is required.

```
SQL> select dvid, vertices.v
       from connectionsGE$, connectionsVT$ vertices
       where svid=1
          and el='collaborates'
          and dvid=vertices.vid
          and vertices.k='occupation';
```

The preceding example's output and execution plan may be as follows.

```
21 67th United States Secretary of State
22 68th United States Secretary of State
23 chancellor
28 7th president of Iran
19 junior United States Senator from New York
...
```

Id	Operation	Bytes	Cost (%CPU)	Time	Pstart	Pstop	TQ	IN-OUT	PQ	Distrib	Rows
0	SELECT STATEMENT										7
525	7 (0)	00:00:01									
1	PX COORDINATOR										
2	PX SEND QC (RANDOM)						:TQ10000				7
525	7 (0)	00:00:01					Q1,00	P->S	QC (RAND)		
3	NESTED LOOPS										7
525	7 (0)	00:00:01					Q1,00	PCWP			
4	PX PARTITION HASH ALL										10
250	2 (0)	00:00:01	1	8			Q1,00	PCWC			
* 5	TABLE ACCESS BY LOCAL INDEX ROWID BATCHED								CONNECTIONSGE\$		10
250	2 (0)	00:00:01	1	8			Q1,00	PCWP			
* 6	INDEX RANGE SCAN								CONNECTIONSXSE\$		20
	1 (0)	00:00:01	1	8			Q1,00	PCWP			
7	PARTITION HASH ITERATOR										1
	0 (0)	00:00:01	KEY	KEY			Q1,00	PCWP			
* 8	TABLE ACCESS BY LOCAL INDEX ROWID								CONNECTIONSVT\$		
			KEY	KEY			Q1,00	PCWP			
* 9	INDEX UNIQUE SCAN								CONNECTIONSXQV\$		1
	0 (0)	00:00:01	KEY	KEY			Q1,00	PCWP			

Predicate Information (identified by operation id):

```
5 - filter(INTERNAL_FUNCTION("EL") AND "EL"=U'collaborates')
6 - access("SVID"=1)
```



```

8 - filter(INTERNAL_FUNCTION("VERTICES"."v"))
9 - access("DVID"="VERTICES"."VID" AND "VERTICES"."K"=U'occupation')
   filter(INTERNAL_FUNCTION("VERTICES"."K"))

```

Example 5-13 Find a Person's Enemies and Aggregate Them by Their Country

The following SQL statement finds enemies (that is, those with the `feuds` relationship) of the vertex with ID 1, and aggregates them by their countries. A join with the vertices table (VT\$) is required.

```

SQL> select vertices.v, count(1)
      from connectionsGE$, connectionsVT$ vertices
      where svid=1
         and el='feuds'
         and dvid=vertices.vid
         and vertices.k='country'
      group by vertices.v;

```

The example's output and execution plan may be as follows. In this case, the vertex with ID 1 has 3 enemies in the United States and 1 in Russia.

```

United States    3
Russia           1

```

Id	Operation	Name	Rows
Bytes	Cost (%CPU) Time	TQ IN-OUT PQ Distrib	
0	SELECT STATEMENT		5
375	5 (20) 00:00:01		
1	PX COORDINATOR		
2	PX SEND QC (RANDOM)	:TQ10001	5
375	5 (20) 00:00:01	Q1,01 P->S QC (RAND)	
3	HASH GROUP BY		5
375	5 (20) 00:00:01	Q1,01 PCWP	
4	PX RECEIVE		5
375	5 (20) 00:00:01	Q1,01 PCWP	
5	PX SEND HASH	:TQ10000	5
375	5 (20) 00:00:01	Q1,00 P->P HASH	
6	HASH GROUP BY		5
375	5 (20) 00:00:01	Q1,00 PCWP	
7	NESTED LOOPS		5
375	4 (0) 00:00:01	Q1,00 PCWP	
8	PX PARTITION HASH ALL		5
125	2 (0) 00:00:01	1 8 Q1,00 PCWC	
* 9	TABLE ACCESS BY LOCAL INDEX ROWID BATCHED	CONNECTIONSGE\$	5
125	2 (0) 00:00:01	1 8 Q1,00 PCWP	
* 10	INDEX RANGE SCAN	CONNECTIONSXSE\$	20
	1 (0) 00:00:01	1 8 Q1,00 PCWP	
11	PARTITION HASH ITERATOR		1
	0 (0) 00:00:01	KEY KEY Q1,00 PCWP	
* 12	TABLE ACCESS BY LOCAL INDEX ROWID	CONNECTIONSVT\$	
		KEY KEY Q1,00 PCWP	
* 13	INDEX UNIQUE SCAN	CONNECTIONSXQV\$	1
	0 (0) 00:00:01	KEY KEY Q1,00 PCWP	

Predicate Information (identified by operation id):

```

 9 - filter(INTERNAL_FUNCTION("EL") AND "EL"=U'feuds')
10 - access("SVID"=1)
12 - filter(INTERNAL_FUNCTION("VERTICES"."v"))
13 - access("DVID"="VERTICES"."VID" AND "VERTICES"."K"=U'country')
      filter(INTERNAL_FUNCTION("VERTICES"."K"))

```

Example 5-14 Find a Person's Collaborators, and aggregate and sort them

The following SQL statement finds the collaborators of the vertex with ID 1, aggregates them by their country, and sorts them in ascending order.

```

SQL> select vertices.v, count(1)
      from connectionsGE$, connectionsVT$ vertices
      where svid=1
         and el='collaborates'
         and dvid=vertices.vid
         and vertices.k='country'
      group by vertices.v
      order by count(1) asc;

```

The example output and execution plan may be as follows. In this case, the vertex with ID 1 has the most collaborators in the United States.

```

Germany          1
Japan             1
Iran              1
United States    7

```


| Id | Operation | Name |
Rows | Bytes | Cost (%CPU) | Time | Pstart| Pstop | TQ | IN-OUT | PQ
Distrib

| 0 | SELECT STATEMENT | |
10 | 750 | 9 (23) | 00:00:01 | | | | | | |
1	PX COORDINATOR								
2	PX SEND QC (ORDER)	:TQ10002							
10	750	9 (23)	00:00:01			Q1,02	P->S	QC	
(ORDER)									
3	SORT ORDER BY								
10	750	9 (23)	00:00:01			Q1,02	PCWP		
4	PX RECEIVE								
10	750	9 (23)	00:00:01			Q1,02	PCWP		
5	PX SEND RANGE	:TQ10001							
10	750	9 (23)	00:00:01			Q1,01	P->P	RANGE	
6	HASH GROUP BY								

10		750		9	(23)		00:00:01				Q1,01		PCWP				
	7																10
	750		9	(23)		00:00:01					Q1,01		PCWP				
	8												:	TQ10000			10
	750		9	(23)		00:00:01					Q1,00		P->P		HASH		
	9																10
	750		9	(23)		00:00:01					Q1,00		PCWP				
	10																10
	750		7	(0)		00:00:01					Q1,00		PCWP				
	11																10
	250		2	(0)		00:00:01		1		8		Q1,00		PCWC			
*	12																10
	250		2	(0)		00:00:01		1		8		Q1,00		PCWP			
*	13																20
			1	(0)		00:00:01		1		8		Q1,00		PCWP			
	14																1
			0	(0)		00:00:01		KEY		KEY		Q1,00		PCWP			
*	15																
								KEY		KEY		Q1,00		PCWP			
*	16																1
			0	(0)		00:00:01		KEY		KEY		Q1,00		PCWP			

 Predicate Information (identified by operation id):

```

12 - filter(INTERNAL_FUNCTION("EL") AND "EL"=U'collaborates')
13 - access("SVID"=1)
15 - filter(INTERNAL_FUNCTION("VERTICES"."v"))
16 - access("DVID"="VERTICES"."VID" AND "VERTICES"."K"=U'country')
      filter(INTERNAL_FUNCTION("VERTICES"."K"))
  
```

5.7.4 Navigation Options: CONNECT BY and Parallel Recursion

The CONNECT BY clause and parallel recursion provide options for advanced navigation and querying.

- CONNECT BY lets you navigate and find matches in a hierarchical order. To follow outgoing edges, you can use prior dvid = svid to guide the navigation.
- Parallel recursion lets you perform navigation up to a specified number of hops away.

The examples use a property graph named connections.

Example 5-15 CONNECT WITH

The following SQL statement follows the outgoing edges by 1 hop.

```

SQL> select G.dvid
       from connectionsGE$ G
       start with svid = 1
       connect by nocycle prior dvid = svid and level <= 1;
  
```

The preceding example's output and execution plan may be as follows.

```

2
3
4
5
  
```

```

6
7
8
9
10
...
-----
| Id | Operation                               | Name                               | Rows | Bytes | Cost
(%CPU)| Time      | Pstart| Pstop | TQ   | IN-OUT| PQ Distrib |
-----
| 0 | SELECT STATEMENT                         |                                     | 7 | 273 | 3
(67)| 00:00:01 |      |      |      |      |      |
|* 1 | CONNECT BY WITH FILTERING                |                                     |      |      |
| 2 | PX COORDINATOR                           |                                     |      |      |
| 3 | PX SEND QC (RANDOM)                       | :TQ10000                           | 2 | 12 | 0
(0)| 00:00:01 |      |      | Q1,00 | P->S | QC (RAND) |
| 4 | PX PARTITION HASH ALL                     |                                     | 2 | 12 | 0
(0)| 00:00:01 | 1 | 8 | Q1,00 | PCWC |
|* 5 | INDEX RANGE SCAN                         | CONNECTIONSXSE$                     | 2 | 12 | 0
(0)| 00:00:01 | 1 | 8 | Q1,00 | PCWP |
|* 6 | FILTER                                    |                                     |      |      |
| 7 | NESTED LOOPS                             |                                     | 5 | 95 | 1
(0)| 00:00:01 |      |      |      |      |
| 8 | CONNECT BY PUMP                           |                                     |      |      |
| 9 | PARTITION HASH ALL                       |                                     | 2 | 12 | 0
(0)| 00:00:01 | 1 | 8 |      |      |
|* 10 | INDEX RANGE SCAN                        | CONNECTIONSXSE$                     | 2 | 12 | 0
(0)| 00:00:01 | 1 | 8 |      |      |
-----

```

Predicate Information (identified by operation id):

- ```

1 - access("SVID"=PRIOR "DVID")
 filter(LEVEL<=2)
5 - access("SVID"=1)
6 - filter(LEVEL<=2)
10 - access("connect$_by$_pump$_002"."prior dvid "="SVID")

```

To extend from 1 hop to multiple hops, change 1 in the preceding example to another integer. For example, to change it to 2 hops, specify: `level <= 2`

### Example 5-16 Parallel Recursion

The following SQL statement uses recursion within the WITH clause to perform navigation up to 4 hops away, a using recursively defined graph expansion: `g_exp` references `g_exp` in the query, and that defines the recursion. The example also uses the `PARALLEL` optimizer hint for parallel execution.

```

SQL> WITH g_exp(svid, dvid, depth) as
(
 select svid as svid, dvid as dvid, 0 as depth
 from connectionsGE$

```

```

 where svid=1
 union all
 select g2.svid, g1.dvid, g2.depth + 1
 from g_exp g2, connectionsGE$ g1
 where g2.dvid=g1.svid
 and g2.depth <= 3
)
select /*+ parallel(4) */ dvid, depth
 from g_exp
 where svid=1
;

```

The example's output and execution plan may be as follows. Note that `CURSOR DURATION MEMORY` is chosen in the execution, which indicates the graph expansion stores the intermediate data in memory.

```

 22 4
 25 4
 24 4
 1 4

 23 4
 33 4
 22 4
 22 4


```

Execution Plan

| Id   | Operation                                        | Name                       |      |        |       |           |        |           |         |
|------|--------------------------------------------------|----------------------------|------|--------|-------|-----------|--------|-----------|---------|
| Rows | Bytes                                            | Cost (%CPU)                | Time | Pstart | Pstop | TQ        | IN-OUT | PQ        | Distrib |
| 0    | SELECT STATEMENT                                 |                            |      |        |       |           |        |           |         |
| 801  | 31239   147 (0)                                  | 00:00:01                   |      |        |       |           |        |           |         |
| 1    | TEMP TABLE TRANSFORMATION                        |                            |      |        |       |           |        |           |         |
| 2    | LOAD AS SELECT ( <b>CURSOR DURATION MEMORY</b> ) | SYS_TEMP_0FD9D6614_11CB2D2 |      |        |       |           |        |           |         |
| 3    | UNION ALL (RECURSIVE WITH) BREADTH FIRST         |                            |      |        |       |           |        |           |         |
| 4    | PX COORDINATOR                                   |                            |      |        |       |           |        |           |         |
| 5    | PX SEND QC (RANDOM)                              | :TQ20000                   |      |        |       |           |        |           |         |
| 2    | 12   0 (0)                                       | 00:00:01                   |      |        |       | Q2,00     | P->S   | QC (RAND) |         |
| 6    | LOAD AS SELECT ( <b>CURSOR DURATION MEMORY</b> ) | SYS_TEMP_0FD9D6614_11CB2D2 |      |        |       |           |        |           |         |
|      |                                                  |                            |      |        |       | Q2,00     | PCWP   |           |         |
| 7    | PX PARTITION HASH ALL                            |                            |      |        |       |           |        |           |         |
| 2    | 12   0 (0)                                       | 00:00:01                   |      | 1      |       | 8   Q2,00 | PCWC   |           |         |
| * 8  | INDEX RANGE SCAN                                 | CONNECTIONSXSE\$           |      |        |       |           |        |           |         |
| 2    | 12   0 (0)                                       | 00:00:01                   |      | 1      |       | 8   Q2,00 | PCWP   |           |         |
| 9    | PX COORDINATOR                                   |                            |      |        |       |           |        |           |         |
| 10   | PX SEND QC (RANDOM)                              | :TQ10000                   |      |        |       |           |        |           |         |
| 799  | 12M   12 (0)                                     | 00:00:01                   |      |        |       | Q1,00     | P->S   | QC (RAND) |         |
| 11   | LOAD AS SELECT ( <b>CURSOR DURATION MEMORY</b> ) | SYS_TEMP_0FD9D6614_11CB2D2 |      |        |       |           |        |           |         |
|      |                                                  |                            |      |        |       | Q1,00     | PCWP   |           |         |

```

|* 12 | HASH JOIN
| | | | | | | | |
| | Q1,00 | PCWP | | | | | |
| 13 | BUFFER SORT (REUSE)
| | | | | | | | |
| | Q1,00 | PCWP | | | | | |
| 14 | PARTITION HASH ALL
| | | | | | | | |
| | 8 | Q1,00 | PCWC | | | | | | |
| 15 | INDEX FAST FULL SCAN
CONNECTIONSXDE$ | 164 | 984 | | | | | |
8 | Q1,00 | PCWP | | | | | | |
| 16 | PX BLOCK ITERATOR
| | | | | | | | |
| | Q1,00 | PCWC | | | | | | |
|* 17 | TABLE ACCESS FULL
SYS_TEMP_0FD9D6614_11CB2D2 | | | | | | | |
| | Q1,00 | PCWP | | | | | | |
| 18 | PX COORDINATOR
| | | | | | | | |
| | | | | | | | |
| 19 | PX SEND QC (RANDOM)
| :TQ30000 | 801 | 31239 | | | | | |
| | Q3,00 | P->S | QC (RAND) | | | | | |
|* 20 | VIEW
| | | | | | | | |
| | Q3,00 | PCWP | | | | | | |
| 21 | PX BLOCK ITERATOR
| | | | | | | | |
| | Q3,00 | PCWC | | | | | | |
| 22 | TABLE ACCESS FULL
SYS_TEMP_0FD9D6614_11CB2D2 | 801 | | 12M | | | | |
| | Q3,00 | PCWP | | | | | | |


```

Predicate Information (identified by operation id):

- 8 - access("SVID"=1)
- 12 - access("G2"."DVID"="G1"."SVID")
- 17 - filter("G2"."INTERNAL\_ITERS\$"=LEVEL AND "G2"."DEPTH"<=3)
- 20 - filter("SVID"=1)

### 5.7.5 Pivot

The PIVOT clause lets you dynamically add columns to a table to create a new table.

The schema design (VT\$ and GE\$) of the property graph is narrow ("skinny") rather than wide ("fat"). This means that if a vertex or edge has multiple properties, those property keys, values, data types, and so on will be stored using multiple rows instead of multiple columns. Such a design is very flexible in the sense that you can add properties dynamically without having to worry about adding too many columns or even reaching the physical maximum limit of number of columns a table may have. However, for some applications you may prefer to have a wide table if the properties are somewhat homogeneous.

**Example 5-17 Pivot**

The following CREATE TABLE ... AS SELECT statement uses PIVOT to add four columns: 'company', 'occupation', 'name', and 'religion'.

```
SQL> CREATE TABLE table pg_wide
as
 with G AS (select vid, k, t, v
 from connectionsVT$
)
 select *
 from G
 pivot (
 min(v) for k in ('company', 'occupation', 'name', 'religion')
);
```

Table created.

The following DESCRIBE statement shows the definition of the new table, including the four added columns. (The output is reformatted for readability.)

```
SQL> DESCRIBE pg_wide;
Name Null? Type

VID NOT NULL NUMBER
T NUMBER (38)
'company'
NVARCHAR2(15000)
'occupation'
NVARCHAR2(15000)
'name'
NVARCHAR2(15000)
'religion'
NVARCHAR2(15000)
```

## 5.7.6 SQL-Based Property Graph Analytics

In addition to the analytical functions offered by the in-memory analyst, the property graph feature in Oracle Spatial and Graph supports several native, SQL-based property graph analytics.

The benefits of SQL-based analytics are:

- Easier analysis of larger graphs that do not fit in physical memory
- Cheaper analysis since no graph data is transferred outside the database
- Better analysis using the current state of a property graph database
- Simpler analysis by eliminating the step of synchronizing an in-memory graph with the latest updates from the graph database

However, when a graph (or a subgraph) fits in memory, then running analytics provided by the in-memory analyst usually provides better performance than using SQL-based analytics.

Because many of the analytics implementation require using intermediate data structures, most SQL- (and PL/SQL-) based analytics APIs have parameters for working tables (wt). A typical flow has the following steps:

1. Prepare the working table or tables.
2. Perform analytics (one or multiple calls).
3. Perform cleanup

The following subtopics provide SQL-based examples of some popular types of property graph analytics.

- [Shortest Path Examples](#)
- [Collaborative Filtering Overview and Examples](#)

### 5.7.6.1 Shortest Path Examples

The following examples demonstrate SQL-based shortest path analytics.

#### Example 5-18 Shortest Path Setup and Computation

Consider shortest path, for example. Internally, Oracle Database uses the bidirectional Dijkstra algorithm. The following code snippet shows an entire prepare, perform, and cleanup workflow.

```

set serveroutput on

DECLARE
 wt1 varchar2(100); -- intermediate working tables
 n number;
 path varchar2(1000);
 weights varchar2(1000);
BEGIN
 -- prepare
 opg_apis.find_sp_prep('connectionsGE$', wt1);
 dbms_output.put_line('working table name ' || wt1);

 -- compute
 opg_apis.find_sp(
 'connectionsGE$',
 1, -- start vertex ID
 53, -- destination vertex ID
 wt1, -- working table (for Dijkstra
expansion)
 dop => 1, -- degree of parallelism
 stats_freq=>1000, -- frequency to collect statistics
 path_output => path, -- shortest path (a sequence of
vertices)
 weights_output => weights, -- edge weights
 options => null
);
 dbms_output.put_line('path ' || path);
 dbms_output.put_line('weights ' || weights);

 -- cleanup (commented out here; see text after the example)
 -- opg_apis.find_sp_cleanup('connectionsGE$', wt1);

```



```
END;
/
```

This example may produce the following output. Note that if **no** working table name is provided, the preparation step will automatically generate a temporary table name and create it. Because the temporary working table name uses the session ID, your output will probably be different.

```
working table name "CONNECTIONSGE$$TWFS12"
path 1 3 52 53
weights 4 3 1 1 1
```

PL/SQL procedure successfully completed.

If you want to know the definition of the working table or tables, then skip the cleanup phase (as shown in the preceding example that comments out the call to `find_sp_cleanup`). After the computation is done, you can describe the working table or tables.

```
SQL> describe "CONNECTIONSGE$$TWFS12"
Name Null? Type

NID NUMBER
D2S NUMBER
P2S NUMBER
D2T NUMBER
P2T NUMBER
F NUMBER(38)
B NUMBER(38)
```

For advanced users who want to try different table creation options, such as using in-memory or advanced compression, you can pre-create the preceding working table and pass the name in.

### Example 5-19 Shortest Path: Create Working Table and Perform Analytics

The following statements show some advanced options, first creating a working table with the same column structure and basic compression enabled, then passing it to the SQL-based computation. The code optimizes the intermediate table for computations with CREATE TABLE compression and in-memory options.

```
create table connections$MY_EXP(
 NID NUMBER,
 D2S NUMBER,
 P2S NUMBER,
 D2T NUMBER,
 P2T NUMBER,
 F NUMBER(38),
 B NUMBER(38)
) compress nologging;

DECLARE
 wt1 varchar2(100) := 'connections$MY_EXP';
 n number;
 path varchar2(1000);
```

```

weights varchar2(1000);
BEGIN
 dbms_output.put_line('working table name ' || wt1);

 -- compute
 opg_apis.find_sp(
 'connectionsGE$',
 1,
 53,
 wt1,
 dop => 1,
 stats_freq=>1000,
 path_output => path,
 weights_output => weights,
 options => null
);
 dbms_output.put_line('path ' || path);
 dbms_output.put_line('weights ' || weights);

 -- cleanup
 -- opg_apis.find_sp_cleanup('connectionsGE$', wt1);
END;
/

```

At the end of the computation, if the working table has not been dropped or truncated, you can check the content of the working table, as follows. Note that the working table structure may vary between releases.

```

SQL> select * from connections$MY_EXP;

```

|    | NID | D2S        | P2S          | D2T | P2T |
|----|-----|------------|--------------|-----|-----|
| F  | B   |            |              |     |     |
|    | 1   | 0          | 1.000E+100   |     |     |
| 1  | -1  |            |              |     |     |
|    | 53  | 1.000E+100 |              | 0   |     |
| -1 | 1   |            |              |     |     |
|    | 54  | 1.000E+100 |              | 1   | 53  |
| -1 | 1   |            |              |     |     |
|    | 52  | 1.000E+100 |              | 1   | 53  |
| -1 | 1   |            |              |     |     |
|    | 5   | 1          | 1 1.000E+100 |     |     |
| 0  | -1  |            |              |     |     |
|    | 26  | 1          | 1 1.000E+100 |     |     |
| 0  | -1  |            |              |     |     |
|    | 8   | 1000       | 1 1.000E+100 |     |     |
| 0  | -1  |            |              |     |     |
|    | 3   | 1          | 1            | 2   | 52  |
| 0  | 0   |            |              |     |     |
|    | 15  | 1          | 1 1.000E+100 |     |     |
| 0  | -1  |            |              |     |     |
|    | 21  | 1          | 1 1.000E+100 |     |     |
| 0  | -1  |            |              |     |     |
|    | 19  | 1          | 1 1.000E+100 |     |     |

```
0 -1
 ...
```

### Example 5-20 Shortest Path: Perform Multiple Calls to Same Graph

To perform multiple calls to the same graph, only *a single call* to the preparation step is needed. The following shows an example of computing shortest path for multiple pairs of vertices in the same graph.

```
DECLARE
 wt1 varchar2(100); -- intermediate working tables
 n number;
 path varchar2(1000);
 weights varchar2(1000);
BEGIN
 -- prepare
 opg_apis.find_sp_prep('connectionsGE$', wt1);
 dbms_output.put_line('working table name ' || wt1);

 -- find shortest path from vertex 1 to vertex 53
 opg_apis.find_sp('connectionsGE$', 1, 53,
 wt1, dop => 1, stats_freq=>1000, path_output => path, weights_output
=> weights, options => null);
 dbms_output.put_line('path ' || path);
 dbms_output.put_line('weights ' || weights);

 -- find shortest path from vertex 2 to vertex 36
 opg_apis.find_sp('connectionsGE$', 2, 36,
 wt1, dop => 1, stats_freq=>1000, path_output => path, weights_output
=> weights, options => null);
 dbms_output.put_line('path ' || path);
 dbms_output.put_line('weights ' || weights);

 -- find shortest path from vertex 30 to vertex 4
 opg_apis.find_sp('connectionsGE$', 30, 4,
 wt1, dop => 1, stats_freq=>1000, path_output => path, weights_output
=> weights, options => null);
 dbms_output.put_line('path ' || path);
 dbms_output.put_line('weights ' || weights);

 -- cleanup
 opg_apis.find_sp_cleanup('connectionsGE$', wt1);
END;
/
```

The example's output may be as follows: three shortest paths have been found for the multiple pairs of vertices provided.

```
working table name "CONNECTIONSGE$$TWFS12"
path 1 3 52 53
weights 4 3 1 1 1
path 2 36
weights 2 1 1
path 30 21 1 4
weights 4 3 1 1 1
```

PL/SQL procedure successfully completed.

## 5.7.6.2 Collaborative Filtering Overview and Examples

**Collaborative filtering**, also referred to as social filtering, filters information by using the recommendations of other people. Collaborative filtering is widely used in systems that recommend purchases based on purchases by others with similar preferences.

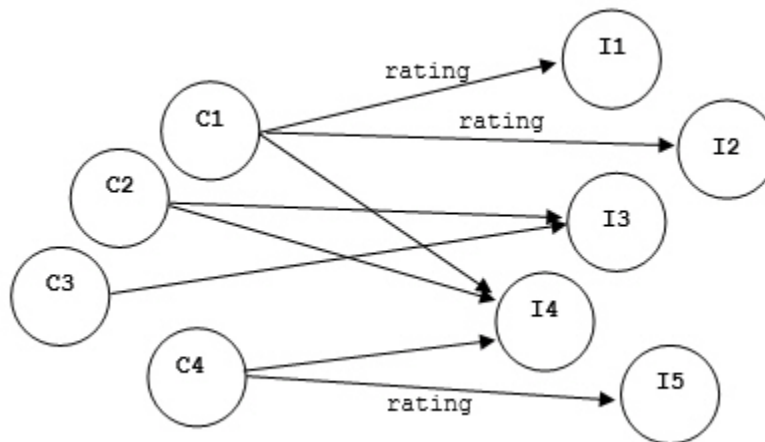
The following examples demonstrate SQL-based collaborative filtering analytics.

### Example 5-21 Collaborative Filtering Setup and Computation

This example shows how to use SQL-based collaborative filtering, specifically using matrix factorization to recommend telephone brands to customers. This example assumes there exists a graph called "PHONES" in the database. This example graph contains customer and item vertices, and edges with a 'rating' label linking some customer vertices to other some item vertices. The rating labels have a numeric value corresponding to the rating that a specific customer (edge OUT vertex) assigned to the specified product (edge IN vertex).

The following figure shows this graph.

**Figure 5-1 Phones Graph for Collaborative Filtering**



```
set serveroutput on
```

```
DECLARE
 wt_l varchar2(32); -- working tables
 wt_r varchar2(32);
 wt_l1 varchar2(32);
 wt_r1 varchar2(32);
 wt_i varchar2(32);
 wt_ld varchar2(32);
 wt_rd varchar2(32);
 edge_tab_name varchar2(32) := 'phonesge$';
 edge_label varchar2(32) := 'rating';
 rating_property varchar2(32) := '';
 iterations integer := 100;
```

```

min_error number := 0.001;
k integer := 5;
learning_rate number := 0.001;
decrease_rate number := 0.95;
regularization number := 0.02;
dop number := 2;
tablespace varchar2(32) := null;
options varchar2(32) := null;
BEGIN

-- prepare
opg_apis.cf_prep(edge_tab_name,wt_l,wt_r,wt_ll,wt_rl,wt_i,wt_ld,wt_rd);
dbms_output.put_line('working table wt_l ' || wt_l);
dbms_output.put_line('working table wt_r ' || wt_r);
dbms_output.put_line('working table wt_ll ' || wt_ll);
dbms_output.put_line('working table wt_rl ' || wt_rl);
dbms_output.put_line('working table wt_i ' || wt_i);
dbms_output.put_line('working table wt_ld ' || wt_ld);
dbms_output.put_line('working table wt_rd ' || wt_rd);

-- compute
opg_apis.cf(edge_tab_name,edge_label,rating_property,iterations,
 min_error,k,learning_rate,decrease_rate,regularization,dop,
 wt_l,wt_r,wt_ll,wt_rl,wt_i,wt_ld,wt_rd,tablespace,options);
END;
/

```

**no**

```

working table wt_l "PHONESGE$$CFL57"
working table wt_r "PHONESGE$$CFR57"
working table wt_ll "PHONESGE$$CFL157"
working table wt_rl "PHONESGE$$CFR157"
working table wt_i "PHONESGE$$CFI57"
working table wt_ld "PHONESGE$$CFLD57"
working table wt_rd "PHONESGE$$CFRD57"

```

PL/SQL procedure successfully completed.

**Example 5-22 Collaborative Filtering: Validating the Intermediate Error**

At the end of every computation, you can check the current error of the algorithm with the following query as long as the data in the working tables has not been already deleted. The following SQL query illustrates how to get the intermediate error of a current run of the collaborative filtering algorithm.

```

SELECT /*+ parallel(48) */ SQRT(SUM((w1-w2)*(w1-w2) +
 <regularization>/2 * (err_reg_l+err_reg_r)) AS err
FROM <wt_i>;

```

Note that the regularization parameter and the working table name (parameter `wt_i`) should be replaced according to the values used when running the `OPG_APIS.CF` algorithm. In the

preceding previous example, replace `<regularization>` with `0.02` and `<wt_i>` with `"PHONESGE$$CFI149"` as follows:

```
SELECT /*+ parallel(48) */ SQRT(SUM((w1-w2)*(w1-w2) + 0.02/2 *
(err_reg_l+err_reg_r))) AS err
FROM "PHONESGE$$CFI149";
```

This query may produce the following output.

```

ERR

4.82163662
```

If the value of the current error is too high or if the predictions obtained from the matrix factorization results of the collaborative filtering are not yet useful, you can run more iterations of the algorithm, by reusing the working tables and the progress made so far. The following example shows how to make predictions using the SQL-based collaborative filtering.

### Example 5-23 Collaborative Filtering: Making Predictions

The result of the collaborative filtering algorithm is stored in the tables `wt_l` and `wt_r`, which are the two factors of a matrix product. These matrix factors should be used when making the predictions of the collaborative filtering.

In a typical flow of the algorithm, the two matrix factors can be used to make the predictions before calling the `OPG_APIS.CF_CLEANUP` procedure, or they can be copied and persisted into other tables for later use. The following example demonstrates the latter case:

```
DECLARE
 wt_l varchar2(32); -- working tables
 wt_r varchar2(32);
 wt_ll varchar2(32);
 wt_rl varchar2(32);
 wt_i varchar2(32);
 wt_ld varchar2(32);
 wt_rd varchar2(32);
 edge_tab_name varchar2(32) := 'phonesge$';
 edge_label varchar2(32) := 'rating';
 rating_property varchar2(32) := '';
 iterations integer := 100;
 min_error number := 0.001;
 k integer := 5;
 learning_rate number := 0.001;
 decrease_rate number := 0.95;
 regularization number := 0.02;
 dop number := 2;
 tablespace varchar2(32) := null;
 options varchar2(32) := null;
BEGIN

 -- prepare

 opg_apis.cf_prep(edge_tab_name,wt_l,wt_r,wt_ll,wt_rl,wt_i,wt_ld,wt_rd);
```

```

-- compute
opg_apis.cf(edge_tab_name,edge_label,rating_property,iterations,
 min_error,k,learning_rate,decrease_rate,regularization,dop,
 wt_l,wt_r,wt_ll,wt_rl,wt_i,wt_ld,wt_rd,tablespace,options);

-- save only these two tables for later predictions
EXECUTE IMMEDIATE 'CREATE TABLE customer_mat AS SELECT * FROM ' || wt_l;
EXECUTE IMMEDIATE 'CREATE TABLE item_mat AS SELECT * FROM ' || wt_r;

-- cleanup
opg_apis.cf_cleanup('phonesge$',wt_l,wt_r,wt_ll,wt_rl,wt_i,wt_ld,wt_rd);
END;
/

```

This example will produce the only the following output.

```
PL/SQL procedure successfully completed.
```

Now that the matrix factors are saved in the tables `customer_mat` and `item_mat`, you can use the following query to check the "error" (difference) between the real values (those values that previously existed in the graph as 'ratings') and the estimated predictions (the result of the matrix multiplication in a certain customer row and item column).

Note that the following query is customized with a join on the vertex table in order return an NVARCHAR property of the vertices (for example, the name property) instead of a numeric ID. This query will return all the predictions for every single customer vertex to every item vertex in the graph.

```

SELECT /*+ parallel(48) */ MIN(vertex1.v) AS customer,
 MIN(vertex2.v) AS item,
 MIN(edges.vn) AS real,
 SUM(l.v * r.v) AS predicted
FROM PHONESGE$ edges,
 CUSTOMER_MAT l,
 ITEM_MAT r,
 PHONESVT$ vertex1,
 PHONESVT$ vertex2
WHERE l.k = r.k
 AND l.c = edges.svid(+)
 AND r.p = edges.dvid(+)
 AND l.c = vertex1.vid
 AND r.p = vertex2.vid
GROUP BY l.c, r.p
ORDER BY l.c, r.p -- This order by clause is optional
;

```

This query may produce an output similar to the following (some rows are omitted for brevity).

| CUSTOMER | ITEM       | REAL | PREDICTED  |
|----------|------------|------|------------|
| Adam     | Apple      | 5    | 3.67375703 |
| Adam     | Blackberry |      | 3.66079652 |
| Adam     | Danger     |      | 2.77049596 |
| Adam     | Ericsson   |      | 4.21764858 |
| Adam     | Figo       |      | 3.10631337 |
| Adam     | Google     | 4    | 4.42429022 |

|      |            |   |            |
|------|------------|---|------------|
| Adam | Huawei     | 3 | 3.4289115  |
| Ben  | Apple      |   | 2.82127589 |
| Ben  | Blackberry | 2 | 2.81132282 |
| Ben  | Danger     | 3 | 2.12761307 |
| Ben  | Ericsson   | 3 | 3.2389595  |
| Ben  | Figo       |   | 2.38550534 |
| Ben  | Google     |   | 3.39765075 |
| Ben  | Huawei     |   | 2.63324582 |
| ...  |            |   |            |
| Don  | Apple      |   | 1.3777496  |
| Don  | Blackberry | 1 | 1.37288909 |
| Don  | Danger     | 1 | 1.03900439 |
| Don  | Ericsson   |   | 1.58172236 |
| Don  | Figo       | 1 | 1.16494421 |
| Don  | Google     |   | 1.65921807 |
| Don  | Huawei     | 1 | 1.28592648 |
| Erik | Apple      | 3 | 2.80809351 |
| Erik | Blackberry | 3 | 2.79818695 |
| Erik | Danger     |   | 2.11767182 |
| Erik | Ericsson   | 3 | 3.2238255  |
| Erik | Figo       |   | 2.3743591  |
| Erik | Google     | 3 | 3.38177526 |
| Erik | Huawei     | 3 | 2.62094201 |

If you want to check only some rows to decide whether the prediction results are ready or more iterations of the algorithm should be run, the previous query can be wrapped in an outer query. The following example will select only the first 11 results.

```

SELECT /*+ parallel(48) */ * FROM (
SELECT /*+ parallel(48) */ MIN(vertex1.v) AS customer,
 MIN(vertex2.v) AS item,
 MIN(edges.vn) AS real,
 SUM(l.v * r.v) AS predicted

FROM PHONESGE$ edges,
 CUSTOMER_MAT l,
 ITEM_MAT r,
 PHONESVT$ vertex1,
 PHONESVT$ vertex2
WHERE l.k = r.k
 AND l.c = edges.svid(+)
 AND r.p = edges.dvid(+)
 AND l.c = vertex1.vid
 AND r.p = vertex2.vid
GROUP BY l.c, r.p
ORDER BY l.c, r.p
) WHERE rownum <= 11;

```

This query may produce an output similar to the following.

| CUSTOMER | ITEM       | REAL | PREDICTED  |
|----------|------------|------|------------|
| Adam     | Apple      | 5    | 3.67375703 |
| Adam     | Blackberry |      | 3.66079652 |
| Adam     | Danger     |      | 2.77049596 |
| Adam     | Ericsson   |      | 4.21764858 |
| Adam     | Figo       |      | 3.10631337 |
| Adam     | Google     | 4    | 4.42429022 |
| Adam     | Huawei     | 3    | 3.4289115  |



|     |            |   |            |
|-----|------------|---|------------|
| Ben | Apple      |   | 2.82127589 |
| Ben | Blackberry | 2 | 2.81132282 |
| Ben | Danger     | 3 | 2.12761307 |
| Ben | Ericsson   | 3 | 3.2389595  |

To get a prediction for a specific vertex (customer, item, or both) the query can be restricted with the desired ID values. For example, to get the predicted value of vertex 1 (customer) and vertex 105 (item), you can use the following query.

```
SELECT /*+ parallel(48) */ MIN(vertex1.v) AS customer,
 MIN(vertex2.v) AS item,
 MIN(edges.vn) AS real,
 SUM(l.v * r.v) AS predicted
FROM PHONESGE$ edges,
 CUSTOMER_MAT l,
 ITEM_MAT r,
 PHONESVT$ vertex1,
 PHONESVT$ vertex2
WHERE l.k = r.k
 AND l.c = edges.svid(+)
 AND r.p = edges.dvid(+)
 AND l.c = vertex1.vid
 AND vertex1.vid = 1 /* Remove to get all predictions for item 105 */
 AND r.p = vertex2.vid
 AND vertex2.vid = 105 /* Remove to get all predictions for customer 1 */
 /* Remove both lines to get all predictions */
GROUP BY l.c, r.p
ORDER BY l.c, r.p;
```

This query may produce an output similar to the following.

| CUSTOMER | ITEM     | REAL | PREDICTED  |
|----------|----------|------|------------|
| Adam     | Ericsson |      | 4.21764858 |

## 5.8 Creating Property Graph Views on an RDF Graph

With Oracle Graph, you can view RDF data as a property graph to execute graph analytics operations by creating property graph views over an RDF graph stored in Oracle Database.

Given an RDF model (or a virtual model), the property graph feature creates two views, a <graph\_name>VT\$ view for vertices and a <graph\_name>GE\$ view for edges.

The `PGUtils.createPropertyGraphViewOnRDF` method lets you customize a property graph view over RDF data:

```
public static void createPropertyGraphViewOnRDF(Connection conn /* a Connection
instance to Oracle database */,
 String pgGraphName /* the name of the property graph to be created */,
 String rdfModelName /* the name of the RDF model */,
 boolean virtualModel /* a flag represents if the RDF model
 is virtual model or not;
 true - virtual mode, false - normal model*/,
 RDFPredicate[] predListForVertexAttrs /* an array of RDFPredicate objects
specifying how to create vertex view using these predicates; each RDFPredicate
includes two fields: an URL of the RDF predicate, the corresponding name of vertex key
in the Property Graph. The mapping from RDF predicates to vertex keys will be created
```

```
based on this parameter. */,
 RDFPredicate[] predListForEdges /* an array of RDFPredicate specifying how
to create edge view using these predicates; each RDFPredicate includes two (or
three) fields: an URL of the RDF predicate, the edge label in the Property
Graph, the weight of the edge (optional). The mapping from RDF predicates to
edges will be created based on this parameter. */)

```

This operation requires the name of the property graph, the name of the RDF Model used to generate the Property Graph view, and a set of mappings determining how triples will be parsed into vertices or edges. The `createPropertyGraphViewOnRDF` method requires a *key/value mapping* array specifying how RDF predicates are mapped to Key/Value properties for vertices, and an *edge mapping* array specifying how RDF predicates are mapped to edges. The `PGUtils.RDFPredicate` API lets you create a map from RDF assertions to vertices/edges.

Vertices are created based on the triples matching at least one of the RDF predicates in the key/value mappings. Each triple satisfying one of the RDF predicates defined in the mapping array is parsed into a vertex with ID based on the internal RDF resource ID of the subject of the triple, and a key/value pair whose key is defined by the mapping itself and whose value is obtained from the object of the triple.

The following example defines a key/value mapping of the RDF predicate URI `http://purl.org/dc/elements/1.1/title` to the key/value property with property name `title`.

```
String titleURL = "http://purl.org/dc/elements/1.1/title";
// create an RDFPredicate to specify how to map the RDF predicate to vertex keys
RDFPredicate titleRDFPredicate
 = RDFPredicate.getInstance(titleURL /* RDF Predicate URI */ ,
 "title" /* property name */);

```

Edges are created based on the triples matching at least one of the RDF predicates in the edge mapping array. Each triple satisfying the RDF predicate defined in the mapping array is parsed into an edge with ID based on the row number, an edge label defined by the mapping itself, a source vertex obtained from the RDF Resource ID of the subject of the triple, and a destination vertex obtained from the RDF Resource ID of the object of the triple. For each triple parsed here, two vertices will be created if they were not generated from the key/value mapping.

The following example defines an edge mapping of the RDF predicate URI `http://purl.org/dc/elements/1.1/reference` to an edge with a label *references* and a weight of 0.5d.

```
String referencesURL = "http://purl.org/dc/terms/references";
// create an RDFPredicate to specify how to map the RDF predicate to edges
RDFPredicate referencesRDFPredicate
 = RDFPredicate.getInstance(referencesURL, "references",
0.5d);

```

The following example creates a property graph view over the RDF model `articles` describing different publications, their authors, and references. The generated property graph will include vertices with some key/value properties that may include `title` and `creator`. The edges in the property graph will be determined by the references among publications.

```
Oracle oracle = null;
Connection conn = null;
OraclePropertyGraph pggraph = null;
try {

```

```

// create the connection instance to Oracle database
OracleDataSource ds = new oracle.jdbc.pool.OracleDataSource();
ds.setURL(jdbcUrl);
conn = (OracleConnection) ds.getConnection(user, password);

// define some string variables for RDF predicates
String titleURL = "http://purl.org/dc/elements/1.1/title";
String creatorURL = "http://purl.org/dc/elements/1.1/creator";
String serialnumberURL = "http://purl.org/dc/elements/1.1/serialnumber";
String widthURL = "http://purl.org/dc/elements/1.1/width";
String weightURL = "http://purl.org/dc/elements/1.1/weight";
String onsaleURL = "http://purl.org/dc/elements/1.1/onsale";
String publicationDateURL = "http://purl.org/dc/elements/1.1/publicationDate";
String publicationTimeURL = "http://purl.org/dc/elements/1.1/publicationTime";
String referencesURL = "http://purl.org/dc/terms/references";

// create RDFPredicate[] predsForVertexAttrs to specify how to map
// RDF predicate to vertex keys
RDFPredicate[] predsForVertexAttrs = new RDFPredicate[8];
predsForVertexAttrs[0] = RDFPredicate.getInstance(titleURL, "title");
predsForVertexAttrs[1] = RDFPredicate.getInstance(creatorURL, "creator");
predsForVertexAttrs[2] = RDFPredicate.getInstance(serialnumberURL,
 "serialnumber");
predsForVertexAttrs[3] = RDFPredicate.getInstance(widthURL, "width");
predsForVertexAttrs[4] = RDFPredicate.getInstance(weightURL, "weight");
predsForVertexAttrs[5] = RDFPredicate.getInstance(onsaleURL, "onsale");
predsForVertexAttrs[6] = RDFPredicate.getInstance(publicationDateURL,
 "publicationDate");
predsForVertexAttrs[7] = RDFPredicate.getInstance(publicationTimeURL,
 "publicationTime");

// create RDFPredicate[] predsForEdges to specify how to map RDF predicates to
// edges
RDFPredicate[] predsForEdges = new RDFPredicate[1];
predsForEdges[0] = RDFPredicate.getInstance(referencesURL, "references", 0.5d);

// create PG view on RDF model
PGUtils.createPropertyGraphViewOnRDF(conn, "articles", "articles", false,
 predsForVertexAttrs, predsForEdges);

// get the Property Graph instance
oracle = new Oracle(jdbcUrl, user, password);
pggraph = OraclePropertyGraph.getInstance(oracle, "articles", 24);

System.err.println("----- Vertices from property graph view -----");
pggraph.getVertices();
System.err.println("----- Edges from property graph view -----");
pggraph.getEdges();
}
finally {
 pggraph.shutdown();
 oracle.dispose();
 conn.close();
}

```

Given the following triples in the `articles` RDF model (11 triples), the output property graph will include two vertices, one for `<http://nature.example.com/Article1>` (`v1`) and another one for `<http://nature.example.com/Article2>` (`v2`). For vertex `v1`, it has eight properties, whose values are the same as their RDF predicates. For example, `v1`'s title is *"All about XYZ"*. Similarly for vertex `v2`, it has two properties: `title` and `creator`. The output property

graph will include a single edge (eid:1) from vertex v1 to vertex v2 with an edge label “references” and a weight of 0.5d.

```
<http://nature.example.com/Article1> <http://purl.org/dc/elements/1.1/title>
"All about XYZ"^^xsd:string.
<http://nature.example.com/Article1> <http://purl.org/dc/elements/1.1/creator>
"Jane Smith"^^xsd:string.
<http://nature.example.com/Article1> <http://purl.org/dc/elements/1.1/
serialnumber> "123456"^^xsd:integer.
<http://nature.example.com/Article1> <http://purl.org/dc/elements/1.1/width>
"10.5"^^xsd:float.
<http://nature.example.com/Article1> <http://purl.org/dc/elements/1.1/weight>
"1.08"^^xsd:double.
<http://nature.example.com/Article1> <http://purl.org/dc/elements/1.1/onsale>
"false"^^xsd:boolean.
<http://nature.example.com/Article1> <http://purl.org/dc/elements/1.1/
publicationDate> "2016-03-08"^^xsd:date)
<http://nature.example.com/Article1> <http://purl.org/dc/elements/1.1/
publicationTime> "2016-03-08T10:10:10"^^xsd:dateTime)
<http://nature.example.com/Article2> <http://purl.org/dc/elements/1.1/title> "A
review of ABC"^^xsd:string.
<http://nature.example.com/Article2> <http://purl.org/dc/elements/1.1/creator>
"Joe Bloggs"^^xsd:string.
<http://nature.example.com/Article1> <http://purl.org/dc/terms/references>
<http://nature.example.com/Article2>.
```

The preceding code will produce an output similar as the following. Note that the internal RDF resource ID values may vary across different Oracle databases.

```
----- Vertices from property graph view -----
Vertex ID 7299961478807817799 {creator:str:Jane Smith, onsale:bol:false,
publicationDate:dat:Mon Mar 07 16:00:00 PST 2016, publicationTime:dat:Tue Mar 08
02:10:10 PST 2016, serialnumber:dbl:123456.0, title:str:All about XYZ,
weight:dbl:1.08, width:flo:10.5}
Vertex ID 7074365724528867041 {creator:str:Joe Bloggs, title:str:A review of ABC}
----- Edges from property graph view -----
Edge ID 1 from Vertex ID 7299961478807817799 {creator:str:Jane Smith,
onsale:bol:false, publicationDate:dat:Mon Mar 07 16:00:00 PST 2016,
publicationTime:dat:Tue Mar 08 02:10:10 PST 2016, serialnumber:dbl:123456.0,
title:str:All about XYZ, weight:dbl:1.08, width:flo:10.5} =[references]=> Vertex
ID 7074365724528867041 {creator:str:Joe Bloggs, title:str:A review of ABC}
edgeKV[{weight:dbl:0.5}]
```

# 6

## Property Graph Query Language (PGQL)

PGQL is a SQL-like query language for property graph data structures that consist of *vertices* that are connected to other vertices by *edges*, each of which can have key-value pairs (properties) associated with them.

The language is based on the concept of *graph pattern matching*, which allows you to specify patterns that are matched against vertices and edges in a data graph.

### Note:

The graph server (PGX) 21.4.0 supports [PGQL 1.4](#) and earlier versions.

The property graph support provides two ways to execute Property Graph Query Language (PGQL) queries through Java APIs:

- Use the `oracle.pgx.api` Java package to query an in-memory snapshot of a graph that has been loaded into the in-memory analyst (PGX), as described in [Using the In-Memory Graph Server \(PGX\)](#).
- Use the `oracle.pg.rdbms.pgql` Java package to directly query graph data stored in Oracle Database, as described in [Executing PGQL Queries Against Property Graph Schema Tables](#).

For more information about PGQL, see <https://pgql-lang.org>.

- [Creating a Property Graph using PGQL](#)
- [Pattern Matching with PGQL](#)
- [Edge Patterns Have a Direction with PGQL](#)
- [Vertex and Edge Labels with PGQL](#)
- [Variable-Length Paths with PGQL](#)
- [Aggregation and Sorting with PGQL](#)
- [Supported PGQL Features and Limitations](#)
- [Executing PGQL Queries Against the In-Memory Graph Server \(PGX\)](#)  
This section describes the Java APIs that are used to execute PGQL queries in the In-Memory graph server (PGX).
- [Executing PGQL Queries Directly Against Oracle Database](#)  
This topic explains how you can execute PGQL queries directly against the graph in Oracle Database (as opposed to in-memory).

### 6.1 Creating a Property Graph using PGQL

`CREATE PROPERTY GRAPH` is a PGQL DDL statement to create a graph from database tables. The graph is stored in the property graph schema.

The CREATE PROPERTY GRAPH statement starts with the name you give the graph, followed by a set of vertex tables and edge tables. The graph can have no vertex tables or edge tables (an empty graph), or vertex tables and no edge tables (a graph with only vertices and no edges), or both vertex tables and edge tables (a graph with vertices and edges). However, a graph cannot be specified with only edge tables and no vertex tables.

Consider the following example:

- **PERSONS** is a table with columns ID, NAME, and ACCOUNT\_NUMBER. A row is added to this table for every person who has an account.
- **TRANSACTIONS** is a table with columns FROM\_ACCOUNT, TO\_ACCOUNT, DATE, and AMOUNT. A row is added into this table in the database every time money is transferred from a FROM\_ACCOUNT to a TO\_ACCOUNT.

A straightforward mapping of tables to graphs is as follows. The graph concepts mapped are: vertices, edges, labels, properties.

- **Vertex tables:** A table that contains data entities is a vertex table.
  - Each row in the vertex table is a vertex.
  - The columns in the vertex table are properties of the vertex.
  - The name of the vertex table is the default label for this set of vertices. Alternatively, you can specify a label name as part of the CREATE PROPERTY GRAPH statement.
- **Edge tables:** An edge table can be any table that links two vertex tables, or a table that has data that indicates an action from a source entity to a target entity. For example, a transfer of money from FROM\_ACCOUNT to TO\_ACCOUNT is a natural edge.
  - Foreign key relationships can give guidance on what links are relevant in your data. CREATE PROPERTY GRAPH will default to using foreign key relationships to identify edges.
  - Some of the properties of an edge table can be the properties of the edge. For example, an edge from FROM\_ACCOUNT to TO\_ACCOUNT can have properties DATE and AMOUNT.
  - The name of an edge table is the default label for this set of edges. Alternatively, you can specify a label name as part of the CREATE PROPERTY GRAPH statement.
- **Keys:**
  - **Keys in a vertex table:** The key of a vertex table identifies a unique vertex in the graph. The key can be specified in the CREATE PROPERTY GRAPH statement; otherwise, it defaults to the primary key of the table. If there are duplicate rows in the table, the CREATE PROPERTY GRAPH statement will return an error.
  - **Key in an edge table:** The key of an edge table uniquely identifies an edge in the graph. The KEY clause when specifying source and destination vertices uniquely identifies the source and destination vertices.

The following is an example CREATE PROPERTY GRAPH statement for the tables PERSONS and TRANSACTIONS.

```
CREATE PROPERTY GRAPH bank_transfers
 VERTEX TABLES (persons KEY(account_number))
 EDGE TABLES (
 transactions KEY (from_acct, to_acct, date, amount)
 SOURCE KEY (from_account) REFERENCES persons
 DESTINATION KEY (to_account) REFERENCES persons
 PROPERTIES (date, amount)
)
```

- **Table aliases:** Vertex and edge tables must have unique names. If you need to identify multiple vertex tables from the same relational table, or multiple edge tables from the same relational table, you must use aliases. For example, you can create two vertex tables PERSONS and PERSONS\_ID from one table PERSONS, as in the following example.

```
CREATE PROPERTY GRAPH bank_transfers
 VERTEX TABLES (persons KEY(account_number)
 persons_id AS persons KEY(id))
```

- **REFERENCES clause:** This connects the source and destination vertices of an edge to the corresponding vertex tables.

For more details, see: <https://pgql-lang.org/spec/latest/#creating-a-property-graph>.

The following table lists the different ways you can create a property graph using the CREATE PROPERTY GRAPH statement:

**Table 6-1 CREATE PROPERTY GRAPH Statement Support**

| Method                                                                                                                      | More Information                                                                  |
|-----------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|
| Create a property graph in the in-memory graph server (PGX) using the <code>oracle.pgx.api</code> Java package              | <a href="#">Java APIs for Executing CREATE PROPERTY GRAPH Statements</a>          |
| Create a property graph in the in-memory graph server (PGX) using the <code>pypgx.api</code> Python package                 | <a href="#">Python APIs for Executing CREATE PROPERTY GRAPH Statements</a>        |
| Create a property graph in Oracle Database (Property Graph Schema) using the <code>oracle.pg.rdbms.pgql</code> Java package | <a href="#">Creating Property Graphs through CREATE PROPERTY GRAPH Statements</a> |
| Create a property graph in Oracle Database (Property Graph Schema) using the <code>opgpy.pgql</code> Python package         | <a href="#">Creating a Property Graph Using the Python Client</a>                 |
| Create a property graph view on Oracle Database tables                                                                      | <a href="#">Creating a Property Graph View</a>                                    |

## 6.2 Pattern Matching with PGQL

Pattern matching is done by specifying one or more path patterns in the MATCH clause. A single path pattern matches a linear path of vertices and edges, while more complex patterns can be matched by combining multiple path patterns, separated by comma. Value

expressions (similar to their SQL equivalents) are specified in the WHERE clause and let you filter out matches, typically by specifying constraints on the properties of the vertices and edges

For example, assume a graph of TCP/IP connections on a computer network, and you want to detect cases where someone logged into one machine, from there into another, and from there into yet another. You would query for that pattern like this:

```
SELECT id(host1) AS id1, id(host2) AS id2, id(host3) AS id3 /*
choose what to return */
FROM MATCH
 (host1) -[connection1]-> (host2) -[connection2]-> (host3) /*
single linear path pattern to match */
WHERE
 connection1.toPort = 22 AND connection1.opened = true AND
 connection2.toPort = 22 AND connection2.opened = true AND
 connection1.bytes > 300 AND /*
meaningful amount of data was exchanged */
 connection2.bytes > 300 AND
 connection1.start < connection2.start AND /*
second connection within time-frame of first */
 connection2.start + connection2.duration < connection1.start +
 connection1.duration
GROUP BY id1, id2, id3 /*
aggregate multiple matching connections */
```

For more examples of pattern matching, see the [relevant section of the PGQL specification](#).

## 6.3 Edge Patterns Have a Direction with PGQL

An edge pattern has a direction, as edges in graphs do. Thus, (a) `<-[]-` (b) specifies a case where *b* has an edge pointing at *a*, whereas (a) `-[]->` (b) looks for an edge in the opposite direction.

The following example finds common friends of April and Chris who are older than both of them.

```
SELECT friend.name, friend.dob
FROM MATCH
 (p1:person) -[:likes]-> (friend) <-[:likes]- (p2:person) /* note the arrow directions below */
WHERE
 p1.name = 'April' AND p2.name = 'Chris' AND
 friend.dob > p1.dob AND friend.dob > p2.dob
ORDER BY friend.dob DESC
```

For more examples of edge patterns, see the relevant section of the PGQL specification [here](#).



## 6.4 Vertex and Edge Labels with PGQL

Labels are a way of attaching type information to edges and nodes in a graph, and can be used in constraints in graphs where not all nodes represent the same thing. For example:

```
SELECT p.name
FROM MATCH (p:person) -[e1:likes]-> (m1:movie),
 MATCH (p) -[e2:likes]-> (m2:movie)
WHERE m1.title = 'Star Wars'
 AND m2.title = 'Avatar'
```

For more examples of label expressions, see the relevant section of the PGQL specification [here](#).

## 6.5 Variable-Length Paths with PGQL

Variable-length path patterns have a quantifier like `*` to match a variable number of vertices and edges. Using a `PATH` macro, you can specify a named path pattern at the start of a query that can be embedded into the `MATCH` clause any number of times, by referencing its name. The following example finds all of the common ancestors of Mario and Luigi.

```
PATH has_parent AS () -[:has_father|has_mother]-> ()
SELECT ancestor.name
FROM MATCH (p1:Person) -/:has_parent*/-> (ancestor:Person)
 , MATCH (p2:Person) -/:has_parent*/-> (ancestor)
WHERE
 p1.name = 'Mario' AND
 p2.name = 'Luigi'
```

The preceding path specification also shows the use of anonymous constraints, because there is no need to define names for intermediate edges or nodes that will not be used in additional constraints or query results. Anonymous elements can have constraints, such as `[:has_father|has_mother]` -- the edge does not get a variable name (because it will not be referenced elsewhere), but it is constrained.

For more examples of variable-length path pattern matching, see the relevant section of the PGQL specification [here](#).

## 6.6 Aggregation and Sorting with PGQL

Like SQL, PGQL has support for the following:

- `GROUP BY` to create groups of solutions
- `MIN`, `MAX`, `SUM`, and `AVG` aggregations
- `ORDER BY` to sort results

And for many other familiar SQL constructs.

For `GROUP BY` and aggregation, see the relevant section of the PGQL specification [here](#). For `ORDER BY`, see the relevant section of the PGQL specification [here](#).

## 6.7 Supported PGQL Features and Limitations

The following table provides the complete list of supported and unsupported PGQL functionalities for:

- PGQL on RDBMS (PG\_SCHEMA option)
- PGQL on RDBMS (PG\_VIEW option)
- PGQL on the graph server (PGX)

**Table 6-2 Supported PGQL Functionalities and Limitations**

| Feature                                      | PGQL on RDBMS (PG_SCHEMA option)                                                                                                                                                                                                                                                            | PGQL on RDBMS (PG_VIEW option)                                                                                                                                                                                                                                                              | PGQL on the Graph Server (PGX)                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CREATE PROPERTY GRAPH                        | Supported                                                                                                                                                                                                                                                                                   | Supported                                                                                                                                                                                                                                                                                   | Supported<br>Limitations: <ul style="list-style-type: none"> <li>• No composite keys for vertices</li> </ul>                                                                                                                                                                                                                                          |
| DROP PROPERTY GRAPH                          | Supported                                                                                                                                                                                                                                                                                   | Supported                                                                                                                                                                                                                                                                                   | Not Supported                                                                                                                                                                                                                                                                                                                                         |
| Fixed-length pattern matching                | Supported                                                                                                                                                                                                                                                                                   | Supported                                                                                                                                                                                                                                                                                   | Supported                                                                                                                                                                                                                                                                                                                                             |
| Variable-length pattern matching goals       | Supported: <ul style="list-style-type: none"> <li>• Reachability</li> </ul> Limitations: <ul style="list-style-type: none"> <li>• ANY</li> <li>• ANY SHORTEST</li> <li>• TOP k SHORTEST</li> <li>• ALL SHORTEST</li> <li>• ANY CHEAPEST</li> <li>• TOP k CHEAPEST</li> <li>• ALL</li> </ul> | Supported: <ul style="list-style-type: none"> <li>• Reachability</li> <li>• ANY</li> <li>• ANY SHORTEST</li> <li>• TOP k SHORTEST</li> </ul> Limitations: <ul style="list-style-type: none"> <li>• ALL SHORTEST</li> <li>• ANY CHEAPEST</li> <li>• TOP k CHEAPEST</li> <li>• ALL</li> </ul> | Supported: <ul style="list-style-type: none"> <li>• Reachability</li> <li>• ANY</li> <li>• ANY SHORTEST</li> <li>• TOP k SHORTEST</li> <li>• ALL SHORTEST</li> <li>• ANY CHEAPEST</li> <li>• TOP k CHEAPEST</li> <li>• ALL</li> </ul>                                                                                                                 |
| Variable-length pattern matching quantifiers | Supported: <ul style="list-style-type: none"> <li>• *</li> <li>• +</li> <li>• ?</li> <li>• { n }</li> <li>• { n, }</li> <li>• { n, m }</li> <li>• { , m }</li> </ul>                                                                                                                        | Supported: <ul style="list-style-type: none"> <li>• *</li> </ul> Limitations: <ul style="list-style-type: none"> <li>• +</li> <li>• ?</li> <li>• { n }</li> <li>• { n, }</li> <li>• { n, m }</li> <li>• { , m }</li> </ul>                                                                  | Supported: <ul style="list-style-type: none"> <li>• *</li> <li>• +</li> <li>• ?</li> <li>• { n }</li> <li>• { n, }</li> <li>• { n, m }</li> <li>• { , m }</li> </ul> Limitations: <ul style="list-style-type: none"> <li>• ? is only supported for reachability</li> <li>• In case of ANY CHEAPEST and TOP k CHEAPEST, only * is supported</li> </ul> |
| GROUP BY                                     | Supported                                                                                                                                                                                                                                                                                   | Supported                                                                                                                                                                                                                                                                                   | Supported                                                                                                                                                                                                                                                                                                                                             |

**Table 6-2 (Cont.) Supported PGQL Functionalities and Limitations**

| Feature                                                                                                                                         | PGQL on RDBMS (PG_SCHEMA option)                                                                                                                                                         | PGQL on RDBMS (PG_VIEW option)                                                                                                                                                   | PGQL on the Graph Server (PGX)                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HAVING                                                                                                                                          | Supported                                                                                                                                                                                | Supported                                                                                                                                                                        | Supported                                                                                                                                                                                                                                                                                                                                                           |
| Aggregations                                                                                                                                    | Supported: <ul style="list-style-type: none"> <li>COUNT</li> <li>MIN, MAX, AVG, SUM</li> </ul> Limitations: <ul style="list-style-type: none"> <li>LISTAGG</li> <li>ARRAY_AGG</li> </ul> | Supported: <ul style="list-style-type: none"> <li>COUNT</li> <li>MIN, MAX, AVG, SUM, LISTAGG</li> </ul> Limitations: <ul style="list-style-type: none"> <li>ARRAY_AGG</li> </ul> | Supported: <ul style="list-style-type: none"> <li>COUNT</li> <li>MIN, MAX, AVG, SUM</li> <li>LISTAGG</li> <li>ARRAY_AGG</li> </ul> Limitations: <ul style="list-style-type: none"> <li>ARRAY_AGG is only supported as <a href="#">horizontal aggregation</a> (in combination with variable-length path) but not in combination with vertical aggregation</li> </ul> |
| DISTINCT <ul style="list-style-type: none"> <li>SELECT DISTINCT</li> <li>Aggregation with DISTINCT (such as, COUNT(DISTINCT e.prop))</li> </ul> | Supported                                                                                                                                                                                | Supported                                                                                                                                                                        | Supported                                                                                                                                                                                                                                                                                                                                                           |
| ORDER BY (+ASC/DESC), LIMIT, OFFSET                                                                                                             | Supported                                                                                                                                                                                | Supported                                                                                                                                                                        | Supported                                                                                                                                                                                                                                                                                                                                                           |
| Data Types                                                                                                                                      | Supported: <ul style="list-style-type: none"> <li>NVARCHAR2 (1500)</li> <li>NUMBER</li> <li>BOOLEAN (stored like NVARCHAR)</li> <li>TIMESTAMP (6) WITH TIME ZONE</li> </ul>              | All available Oracle RDBMS data types supported                                                                                                                                  | Supported: <ul style="list-style-type: none"> <li>INTEGER (32-bit)</li> <li>LONG (64-bit)</li> <li>FLOAT (32-bit)</li> <li>DOUBLE (64-bit)</li> <li>STRING (no maximum length)</li> <li>BOOLEAN</li> <li>DATE</li> <li>TIME</li> <li>TIME WITH TIME ZONE</li> <li>TIMESTAMP</li> <li>TIMESTAMP WITH TIME ZONE</li> </ul>                                            |

**Table 6-2 (Cont.) Supported PGQL Functionalities and Limitations**

| Feature                  | PGQL on RDBMS (PG_SCHEMA option)                                                                                                                                                                                                                                                                                                                                                                                                                                       | PGQL on RDBMS (PG_VIEW option)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | PGQL on the Graph Server (PGX)                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operators                | <p>Supported:</p> <ul style="list-style-type: none"> <li>Relational: +, -, *, /, %, - (unary minus)</li> <li>Arithmetic: =, &lt;&gt;, &lt;, &gt;, &lt;=, &gt;=</li> <li>Logical: AND, OR, NOT</li> </ul> <p>Limitations:</p> <ul style="list-style-type: none"> <li>String:    (concat)</li> </ul>                                                                                                                                                                     | <p>Supported:</p> <ul style="list-style-type: none"> <li>Relational: +, -, *, /, %, - (unary minus)</li> <li>Arithmetic: =, &lt;&gt;, &lt;, &gt;, &lt;=, &gt;=</li> <li>Logical: AND, OR, NOT</li> </ul> <ul style="list-style-type: none"> <li>String:    (concat)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                | <p>Supported:</p> <ul style="list-style-type: none"> <li>Relational: +, -, *, /, %, - (unary minus)</li> <li>Arithmetic: =, &lt;&gt;, &lt;, &gt;, &lt;=, &gt;=</li> <li>Logical: AND, OR, NOT</li> </ul> <ul style="list-style-type: none"> <li>String:    (concat)</li> </ul>                                                                                                                    |
| Functions and predicates | <p>Supported:</p> <ul style="list-style-type: none"> <li>IS NULL, IS NOT NULL</li> <li>JAVA_REGEXPT_LIKE (based on CONTAINS)</li> <li>ABS, CEIL/CEILING, FLOOR, ROUND</li> <li>EXTRACT</li> <li>ID</li> <li>LABEL, HAS_LABEL</li> <li>ALL_DIFFERENT</li> <li>IN_DEGREE, OUT_DEGREE</li> <li>CAST</li> <li>CASE</li> <li>IN and NOT IN</li> </ul> <p>Limitations:</p> <ul style="list-style-type: none"> <li>LOWER, UPPER</li> <li>SUBSTRING</li> <li>LABELS</li> </ul> | <p>Supported are all available functions in the Oracle RDBMS that take the form <code>function_name(arg1, arg2, ...)</code> with optional schema and package qualifiers.</p> <p>Supported PGQL functions/predicates:</p> <ul style="list-style-type: none"> <li>IS NULL, IS NOT NULL</li> <li>JAVA_REGEXPT_LIKE (based on CONTAINS)</li> <li>LOWER, UPPER</li> <li>SUBSTRING</li> <li>ABS, CEIL/CEILING, FLOOR, ROUND</li> <li>EXTRACT</li> <li>ID</li> <li>LABEL, HAS_LABEL</li> <li>ALL_DIFFERENT</li> <li>CAST</li> <li>CASE</li> <li>IN and NOT IN</li> </ul> <p>Limitations:</p> <ul style="list-style-type: none"> <li>LABELS</li> <li>IN_DEGREE, OUT_DEGREE</li> </ul> | <p>Supported:</p> <ul style="list-style-type: none"> <li>IS NULL, IS NOT NULL</li> <li>JAVA_REGEXPT_LIKE (based on CONTAINS)</li> <li>LOWER, UPPER</li> <li>SUBSTRING</li> <li>ABS, CEIL/CEILING, FLOOR, ROUND</li> <li>EXTRACT</li> <li>ID</li> <li>LABEL, LABELS, HAS_LABEL</li> <li>ALL_DIFFERENT</li> <li>IN_DEGREE, OUT_DEGREE</li> <li>CAST</li> <li>CASE</li> <li>IN and NOT IN</li> </ul> |

**Table 6-2 (Cont.) Supported PGQL Functionalities and Limitations**

| Feature                                                                                                                   | PGQL on RDBMS<br>(PG_SCHEMA<br>option) | PGQL on RDBMS<br>(PG_VIEW option)                                                                                                                          | PGQL on the Graph<br>Server (PGX)                                                               |
|---------------------------------------------------------------------------------------------------------------------------|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| User-defined functions                                                                                                    | Not Supported                          | Supported: <ul style="list-style-type: none"> <li>PL/SQL functions</li> <li>Functions created via the Oracle Database Multilingual Engine (MLE)</li> </ul> | Supported: <ul style="list-style-type: none"> <li>Java UDFs</li> <li>JavaScript UDFs</li> </ul> |
| Subqueries: <ul style="list-style-type: none"> <li>Scalar subqueries</li> <li>EXISTS and NOT EXISTS subqueries</li> </ul> | Supported                              | Not Supported                                                                                                                                              | Supported                                                                                       |
| INSERT/UPDATE/<br>DELETE                                                                                                  | Supported                              | Not Supported                                                                                                                                              | Supported                                                                                       |

## 6.8 Executing PGQL Queries Against the In-Memory Graph Server (PGX)

This section describes the Java APIs that are used to execute PGQL queries in the In-Memory graph server (PGX).

- [Getting Started with PGQL](#)
- [PGQL Features Supported in the Graph Server \(PGX\)](#)
- [Java APIs for Executing CREATE PROPERTY GRAPH Statements](#)
- [Python APIs for Executing CREATE PROPERTY GRAPH Statements](#)
- [Java APIs for Executing SELECT Queries](#)  
This section describes the APIs to execute `SELECT` queries in the In-Memory graph server (PGX).
- [Java APIs for Executing UPDATE Queries](#)  
The `UPDATE` queries make changes to existing graphs using the `INSERT`, `UPDATE`, and `DELETE` operations as detailed in the section Graph Modification of the PGQL 1.3 specification.
- [PGQL Queries with Partitioned IDs](#)
- [Security Tools for Executing PGQL Queries](#)  
To safeguard against query injection, bind variables can be used in place of literals while `printIdentifier(String identifier)` can be used in place of identifiers like graph names, labels, and property names.
- [Best Practices for Tuning PGQL Queries](#)  
This section describes best practices regarding memory allocation, parallelism, and query planning.

## 6.8.1 Getting Started with PGQL

This section provides an example on how to get started with PGQL. It assumes a database realm that has been previously set up (follow the steps in [Prepare the Graph Server for Database Authentication](#)). It also assumes that the user has read access to the HR schema.

First, create a graph with employees, departments, and employee works at department, by executing a CREATE PROPERTY GRAPH statement.

### Example 6-1 Creating a graph in the in-memory graph server (PGX)

The following statement creates a graph in the in-memory graph server (PGX)

```
String statement =
 "CREATE PROPERTY GRAPH hr_simplified "
 + " VERTEX TABLES ("
 + " hr.employees LABEL employee "
 + " PROPERTIES ARE ALL COLUMNS EXCEPT (job_id, manager_id,
department_id), "
 + " hr.departments LABEL department "
 + " PROPERTIES (department_id, department_name) "
 + ") "
 + " EDGE TABLES ("
 + " hr.employees AS works_at "
 + " SOURCE KEY (employee_id) REFERENCES employees "
 + " DESTINATION departments "
 + " PROPERTIES (employee_id) "
 + ");";
session.executePgql(statement);

/**
 * To get a handle to the graph, execute:
 */
PgxGraph g = session.getGraph("HR_SIMPLIFIED");

/**
 * You can use this handle to run PGQL queries on this graph.
 * For example, to find the department that "Nandita Sarchand" works
for, execute:
 */
String query =
 "SELECT dep.department_name "
 + "FROM MATCH (emp:Employee) -[:works_at]-> (dep:Department) "
 + "WHERE emp.first_name = 'Nandita' AND emp.last_name = 'Sarchand' "
 + "ORDER BY 1";
PgqlResultSet resultSet = g.queryPgql(query);
resultSet.print();
+-----+
| department_name |
+-----+
| Shipping |
+-----+
```

```

/**
 * To get an overview of the types of vertices and their frequencies,
 * execute:
 */
String query =
 "SELECT label(n), COUNT(*) "
 + "FROM MATCH (n) "
 + "GROUP BY label(n) "
 + "ORDER BY COUNT(*) DESC";
PgqlResultSet resultSet = g.queryPgql(query);
resultSet.print();

+-----+
| label(n) | COUNT(*) |
+-----+
| EMPLOYEE | 107 |
| DEPARTMENT | 27 |
+-----+

/**
 *To get an overview of the types of edges and their frequencies, execute:
 */
String query =
 "SELECT label(n) AS srcLbl, label(e) AS edgeLbl, label(m) AS dstLbl,
COUNT(*) "
 + "FROM MATCH (n) -[e]-> (m) "
 + "GROUP BY srcLbl, edgeLbl, dstLbl "
 + "ORDER BY COUNT(*) DESC";
PgqlResultSet resultSet = g.queryPgql(query);
resultSet.print();

+-----+
| srcLbl | edgeLbl | dstLbl | COUNT(*) |
+-----+
| EMPLOYEE | WORKS_AT | DEPARTMENT | 106 |
+-----+

```

## 6.8.2 PGQL Features Supported in the Graph Server (PGX)

You can view the complete list of supported and unsupported PGQL features by the graph server (PGX) in [Supported PGQL Features and Limitations](#).

However, the following explains certain limitations for a few PGQL features:

- [Limitations on Quantifiers](#)  
Although all quantifiers such as `*`, `+`, and `{1, 4}` are supported for reachability and shortest path patterns, the only quantifier that is supported for cheapest path patterns is `*` (zero or more).
- [Limitations on WHERE and COST Clauses in Quantified Patterns](#)

### 6.8.2.1 Limitations on Quantifiers

Although all quantifiers such as `*`, `+`, and `{1,4}` are supported for reachability and shortest path patterns, the only quantifier that is supported for cheapest path patterns is `*` (zero or more).

### 6.8.2.2 Limitations on WHERE and COST Clauses in Quantified Patterns

The `WHERE` and `COST` clauses in quantified patterns, such as reachability patterns or shortest and cheapest path patterns, are limited to referencing a single variable only.

The following are examples of queries that are not supported because the `WHERE` or `COST` clauses reference two variables `e` and `x` instead of zero or one:

```
... PATH p AS (n) -[e]-> (m) WHERE e.prop > m.prop ...
... SHORTEST ((n) (-[e]-> (x) WHERE e.prop + x.prop > 10)* (m)) ...
... CHEAPEST ((n) (-[e]-> (x) COST e.prop + x.prop) * (m)) ...
```

The following query is supported because the subquery only references a single variable `a` from the outer scope, while the variable `c` does not count since it is newly introduced in the subquery:

```
... PATH p AS (a) -> (b)
 WHERE EXISTS (SELECT * FROM MATCH (a) -> (c)) ...
```

## 6.8.3 Java APIs for Executing CREATE PROPERTY GRAPH Statements

The easiest way to execute a `CREATE PROPERTY GRAPH` statement is through the `PgxSession.executePgql(String statement)` method.

### Example 6-2 Executing a CREATE PROPERTY GRAPH statement

```
String statement =
 "CREATE PROPERTY GRAPH hr_simplified "
 + " VERTEX TABLES ("
 + " hr.employees LABEL employee "
 + " PROPERTIES ARE ALL COLUMNS EXCEPT (job_id, manager_id,
department_id), "
 + " hr.departments LABEL department "
 + " PROPERTIES (department_id, department_name) "
 + ") "
 + " EDGE TABLES ("
 + " hr.employees AS works_at "
 + " SOURCE KEY (employee_id) REFERENCES employees "
 + " DESTINATION departments "
 + " PROPERTIES (employee_id) "
 + ");";
session.executePgql(statement);
PgxGraph g = session.getGraph("HR_SIMPLIFIED");
```



```

/**
 * Alternatively, one can use the prepared statement API, for example:
 */

PgxPreparedStatement stmt = session.preparePgql(statement);
stmt.execute();
stmt.close();
PgxGraph g = session.getGraph("HR_SIMPLIFIED");

```

## 6.8.4 Python APIs for Executing CREATE PROPERTY GRAPH Statements

You can create a property graph by executing the `CREATE PROPERTY GRAPH` statement through the Python API.

### Creating a Property Graph Using the Python Client

- Launch the Python client:

```
./bin/opg4py --base_url https://localhost:7007 --user customer_360
```

- Define and execute the `CREATE PROPERTY GRAPH` statement as shown:

```

statement = (
 "CREATE PROPERTY GRAPH "+ "<graph_name>" + " " +
 "VERTEX TABLES (" +
 "bank_accounts " +
 "KEY(acct_id) " +
 "LABEL Account PROPERTIES (acct_id) " +
 ") " +
 "EDGE TABLES (" +
 "bank_txns " +
 "KEY (txn_id) " +
 "SOURCE KEY (from_acct_id) REFERENCES bank_accounts " +
 "DESTINATION KEY (to_acct_id) REFERENCES bank_accounts " +
 "LABEL Transfer PROPERTIES(amount) " +
 ")")
>>> session.prepare_pgql(statement).execute()

```

where `<graph_name>` is the name of the graph.

The graph gets created and you can verify through the `get_graph` method:

```

>>> graph = session.get_graph("<graph_name>")
>>> graph
PgxGraph(name:<graph_variable>, v: 1000, e: 5001, directed: True,
memory(Mb): 0)

```

## 6.8.5 Java APIs for Executing SELECT Queries

This section describes the APIs to execute `SELECT` queries in the In-Memory graph server (PGX).

- [Executing SELECT Queries Against a Graph in the In-memory Graph Server \(PGX\)](#)  
The `PgxGraph.queryPgql(String query)` method executes the query in the current session. The method returns a `PgqlResultSet`.
- [Executing SELECT Queries Against a PGX Session](#)  
The `PgxSession.queryPgql(String query)` method executes the given query in the session and returns a `PgqlResultSet`.
- [Iterating Through a Result Set](#)  
There are two ways to iterate through a result set: in a JDBC-like manner or using the Java Iterator interface.
- [Printing a Result Set](#)  
The following methods of `PgqlResultSet` (package `oracle.pgx.api`) are used to print a result set:

### 6.8.5.1 Executing SELECT Queries Against a Graph in the In-memory Graph Server (PGX)

The `PgxGraph.queryPgql(String query)` method executes the query in the current session. The method returns a `PgqlResultSet`.

The `ON` clauses inside the `MATCH` clauses can be omitted since the query is executed directly against a PGX graph. For the same reason, the `INTO` clauses inside the `INSERT` clauses can be omitted. However, if you want to explicitly specify graph names in the `ON` and `INTO` clauses, then those graph names have to match the actual name of the graph (`PgxGraph.getName()`).

### 6.8.5.2 Executing SELECT Queries Against a PGX Session

The `PgxSession.queryPgql(String query)` method executes the given query in the session and returns a `PgqlResultSet`.

The `ON` clauses inside the `MATCH` clauses, and the `INTO` clauses inside the `INSERT` clauses, must be specified and cannot be omitted. At this moment, all the `ON` and `INTO` clauses of a query need to reference the same graph since joining data from multiple graphs in a single query is not yet supported.

### 6.8.5.3 Iterating Through a Result Set

There are two ways to iterate through a result set: in a JDBC-like manner or using the Java Iterator interface.

For JDBC-like iterations, the methods in `PgqlResultSet` (package `oracle.pgx.api`) are similar to the ones in `java.sql.ResultSet`. A noteworthy difference is that PGQL's result set interface is based on the new date and time library that was introduced in Java 8, while `java.sql.ResultSet` is based on the legacy `java.util.Date`. To bridge the gap, PGQL's result set provides `getLegacyDate(..)` for applications that still use `java.util.Date`.

A `PgqlResultSet` has a `cursor` that is initially set before the first row. Then, the following methods are available to reposition the cursor:

- `next()` : boolean

- `previous() : boolean`
- `beforeFirst()`
- `afterLast()`
- `first() : boolean`
- `last() : boolean`
- `absolute(long row) : boolean`
- `relative(long rows) : boolean`

Above methods are documented in more detail [here](#).

After the cursor is positioned at the desired row, the following getters are used to obtain values:

- `getObject(int columnIndex) : Object`
- `getObject(String columnName) : Object`
- `getString(int columnIndex) : String`
- `getString(String columnName) : String`
- `getInteger(int columnIndex) : Integer`
- `getInteger(String columnName) : Integer`
- `getLong(int columnIndex) : Long`
- `getLong(String columnName) : Long`
- `getFloat(int columnIndex) : Float`
- `getFloat(String columnName) : Float`
- `getDouble(int columnIndex) : Double`
- `getDouble(String columnName) : Double`
- `getBoolean(int columnIndex) : Boolean`
- `getBoolean(String columnName) : Boolean`
- `getVertexLabels(int columnIndex) : Set<String>`
- `getVertexLabels(String columnName) : Set<String>`
- `getDate(int columnIndex) : LocalDate`
- `getDate(String columnName) : LocalDate`
- `getTime(int columnIndex) : LocalTime`
- `getTime(String columnName) : LocalTime`
- `getTimestamp(int columnIndex) : LocalDateTime`
- `getTimestamp(String columnName) : LocalDateTime`
- `getTimeWithTimezone(int columnIndex) : OffsetTime`
- `getTimeWithTimezone(String columnName) : OffsetTime`
- `getTimestampWithTimezone(int columnIndex) : OffsetDateTime`
- `getTimestampWithTimezone(String columnName) : OffsetDateTime`

- `getLegacyDate(int columnIndex) : java.util.Date`
- `getLegacyDate(String columnName) : java.util.Date`
- `getVertex(int columnIndex) : PgxVertex<ID>`
- `getVertex(String columnName) : PgxVertex<ID>`
- `getEdge(int columnIndex) : PgxEdge`
- `getEdge(String columnName) : PgxEdge`

Above methods are documented in more detail [here](#).

Finally, there is a `PgqlResultSet.close()` which releases the result set's resources, and there is a `PgqlResultSet.getMetaData()` through which the column names and column count can be retrieved.

An example for result set iteration is as follows:

```
PgqlResultSet resultSet = g.queryPgql(
 " SELECT owner.name AS account_holder, SUM(t.amount) AS
total_transacted_with_Nikita "
 + " FROM MATCH (p:Person) -[:ownerOf]-> (account1:Account) "
 + " , MATCH (account1) -[t:transaction]- (account2) "
 + " , MATCH (account2:Account) <-[:ownerOf]- (owner:Person|
Company) "
 + " WHERE p.name = 'Nikita' "
 + " GROUP BY owner");

while (resultSet.next()) {
 String accountHolder = resultSet.getString(1);
 long totalTransacted = resultSet.getLong(2);
 System.out.println(accountHolder + ": " + totalTransacted);
}

resultSet.close();
```

The output of the above example will look like:

```
Oracle: 4501
Camille: 1000
```

In addition, the `PgqlResultSet` is also iterable via the Java Iterator interface. An example of a “for each loop” over the result set is as follows:

```
for (PgxResult result : resultSet) {
 String accountHolder = result.getString(1);
 long totalTransacted = result.getLong(2);
 System.out.println(accountHolder + ": " + totalTransacted);
}
```

The output of the above example will look like:

```
Oracle: 4501
Camille: 1000
```

Note that the same getters that are available for `PgqlResultSet` are also available for `PgxResult`.

### 6.8.5.4 Printing a Result Set

The following methods of `PgqlResultSet` (package `oracle.pgx.api`) are used to print a result set:

- `print() : PgqlResultSet`
- `print(long numResults) : PgqlResultSet`
- `print(long numResults, int from) : PgqlResultSet`
- `print(PrintStream printStream, long numResults, int from) : PgqlResultSet`

For example:

```
g.queryPgql("SELECT COUNT(*) AS numPersons FROM MATCH
(n:Person)").print().close()
+-----+
| numPersons |
+-----+
| 3 |
+-----+
```

Another example:

```
PgqlResultSet resultSet = g.queryPgql(
 " SELECT owner.name AS account_holder, SUM(t.amount) AS
total_transacted_with_Nikita "
 + " FROM MATCH (p:Person) -[:ownerOf]-> (account1:Account) "
 + " , MATCH (account1) -[t:transaction]- (account2) "
 + " , MATCH (account2:Account) <-[:ownerOf]- (owner:Person|Company)
"
 + " WHERE p.name = 'Nikita' "
 + " GROUP BY owner")

resultSet.print().close()
+-----+
| account_holder | total_transacted_with_Nikita |
+-----+
| Camille | 1000.0 |
| Oracle | 4501.0 |
+-----+
```

## 6.8.6 Java APIs for Executing UPDATE Queries

The `UPDATE` queries make changes to existing graphs using the `INSERT`, `UPDATE`, and `DELETE` operations as detailed in the section Graph Modification of the PGQL 1.3 specification.

Note that `INSERT` allows you to insert new vertices and edges into a graph, `UPDATE` allows you to update existing vertices and edges by setting their properties to new values, and `DELETE` allows you to delete vertices and edges from a graph.

- [Updatability of Graphs Through PGQL](#)  
Graph data that is loaded from the Oracle RDBMS or from CSV files into the PGX is not updatable through PGQL right away.
- [Executing UPDATE Queries against a Graph in the in-memory Graph Server \(PGX\)](#)  
To execute `UPDATE` queries against a graph, use the `PgxGraph.executePgql(String query)` method.
- [Executing UPDATE Queries Against a PGX Session](#)  
For now, there is no support for executing `UPDATE` queries against a `PgxSession` and therefore, updates always have to be executed against a `PgxGraph`. To obtain a graph from a session, use the `PgxSession.getGraph(String graphName)` method.
- [Altering the Underlying Schema of a Graph](#)  
The `INSERT` operations can only insert vertices and edges with known labels and properties. Similarly, `UPDATE` operations can only set values of known properties. Thus, new data must always conform to the existing schema of the graph.

### 6.8.6.1 Updatability of Graphs Through PGQL

Graph data that is loaded from the Oracle RDBMS or from CSV files into the PGX is not updatable through PGQL right away.

First, you need to create a copy of the data through the `PgxGraph.clone()` method. The resulting graph is fully updatable.

Consider the following example:

```
// load a graph from the RDBMS or from CSV
PgxGraph g1 = session.readGraphWithProperties("path/to/
graph_config.json");

// create an updatable copy of the graph
PgxGraph g2 = g1.clone("new_graph_name");

// insert an additional vertex into the graph
g2.executePgql("INSERT VERTEX v " +
 " LABELS (Person) " +
 " PROPERTIES (v.firstName = 'Camille', " +
 " v.lastName = ' Mullins'");
```

Additionally, there is also a `PgxGraph.cloneAndExecutePgql(String query, String graphName)` method that combines the last two steps from above example into a single step:

```
// create an updatable copy of the graph while inserting a new vertex
PgxGraph g2_copy = g1.cloneAndExecutePgql(
 "INSERT VERTEX v " +
 " LABELS (Person) " +
 " PROPERTIES (v.firstName = 'Camille', " +
 " v.lastName = ' Mullins') "
 , "new_graph_name");
```

Note that graphs that are created through `PgxGraph.clone()` are local to the session. However, they can be shared with other sessions through the `PgxGraph.publish(..)` methods but then they are no longer updatable through PGQL. Only session-local graphs are updatable but persistent graphs are not.

### 6.8.6.2 Executing UPDATE Queries against a Graph in the in-memory Graph Server (PGX)

To execute UPDATE queries against a graph, use the `PgxGraph.executePgql(String query)` method.

The following is an example of INSERT query:

```
g.executePgql("INSERT VERTEX v " +
 " LABELS (Person) " +
 " PROPERTIES (v.firstName = 'Camille', " +
 " v.lastName = ' Mullins') ");
```

Note that the INTO clause of the INSERT can be omitted. If you use an INTO clause, the graph name in the INTO clause must correspond to the name of the PGX graph (`PgxGraph.getName()`) that the query is executed against.

The following is an example of UPDATE query:

```
// set the date of birth of Camille to 2014-11-15
g.executePgql("UPDATE v SET (v.dob = DATE '2014-11-14') " +
 "FROM MATCH (v:Person) " +
 "WHERE v.firstName = 'Camille' AND v.lastName = ' Mullins' ");
```

The following is an example of DELETE query:

```
// delete Camille from the graph
g.executePgql("DELETE v " +
 "FROM MATCH (v:Person) " +
 "WHERE v.firstName = 'Camille' AND v.lastName = 'Mullins' ");
```

### 6.8.6.3 Executing UPDATE Queries Against a PGX Session

For now, there is no support for executing `UPDATE` queries against a `PgxSession` and therefore, updates always have to be executed against a `PgxGraph`. To obtain a graph from a session, use the `PgxSession.getGraph(String graphName)` method.

### 6.8.6.4 Altering the Underlying Schema of a Graph

The `INSERT` operations can only insert vertices and edges with known labels and properties. Similarly, `UPDATE` operations can only set values of known properties. Thus, new data must always conform to the existing schema of the graph.

However, some PGX APIs exist for updating the schema of a graph: while no APIs exist for adding new labels, new properties can be added through the `PgxGraph.createVertexProperty(PropertyType type, String name)` and `PgxGraph.createEdgeProperty(PropertyType type, String name)` methods. The new properties are attached to each vertex/edge in the graph, irrespective of their labels. Initially the properties are assigned a default value but then the values can be updated through the `UPDATE` statements.

Consider the following example:

```
// load a graph from the RDBMS or from CSV
PgxGraph g = session.readGraphWithProperties("path/to/
graph_config.json");

// add a new property to the graph
g.createVertexProperty(PropertyType.LOCAL_DATE, "dob");

// set the date of birth of Camille to 2014-11-15
g.executePgql("UPDATE v SET (v.dob = DATE '2014-11-14') " +
 "FROM MATCH (v:Person) " +
 "WHERE v.firstName = 'Camille' AND v.lastName = '
Mullins' ");
```

## 6.8.7 PGQL Queries with Partitioned IDs

You can retrieve partitioned IDs using the `id()` function in PGQL.

### PGQL SELECT Queries

The following are a few examples to retrieve partitioned IDs using PGQL `SELECT` queries:

```
g.queryPgql("SELECT id(n) FROM MATCH(n)").print().close()
```

This prints an output similar to:

```
+-----+
| id(n) |
+-----+
| Accounts(2) |
```



```
| Accounts(4) |
| Accounts(6) |
+-----+
```

```
g.queryPgql("SELECT n.name FROM MATCH(n) WHERE id(n) =
'Accounts(1) ").print().close()
```

The output is printed as shown:

```
+-----+
| name |
+-----+
| User1 |
+-----+
```

```
g.queryPgql("SELECT LABEL(n), n.name from MATCH(n) WHERE n.id =
1").print().close()
```

The output is printed as shown:

```
+-----+
| label(n) | name |
+-----+
| Accounts | User1 |
+-----+
```

PGX automatically creates a unique index for keys so that queries with predicates such as `WHERE id(n) = 'Accounts(1)'` and `WHERE n.id = 1` can be efficiently processed by retrieving the vertex in constant time.

### Using Bind Variables

Partitioned IDs can also be passed as bind values into a `PgxPreparedStatement`.

For example:

```
PgxPreparedStatement statement = g.preparePgql("SELECT n.name FROM MATCH (n)
WHERE id(n)= ?")
statement.setString(1, "Accounts(1)")
statement.executeQuery().print().close()
```

This prints the output as shown:

```
+-----+
| name |
+-----+
| User1 |
+-----+
```

### PGQL INSERT Queries

In `INSERT` queries, you must provide a value for the key property if a key property exists. The value is then used for the vertex or edge key.

For example you can execute an `INSERT` as shown:

```
g.executePgql("INSERT VERTEX v LABELS (Accounts) PROPERTIES (v.id = 1001, v.name = 'User1001')")
```

The inserted values can be verified as shown:

```
g.queryPgql("SELECT id(n), n.name FROM MATCH(n) WHERE n.id = 1001").print().close()
```

This prints the output:

```
+-----+
| id(n) | name |
+-----+
| Accounts(1001) | User1001 |
+-----+
```

## 6.8.8 Security Tools for Executing PGQL Queries

To safeguard against query injection, bind variables can be used in place of literals while `printIdentifier(String identifier)` can be used in place of identifiers like graph names, labels, and property names.

- [Using Bind Variables](#)

There are two reasons for using bind variables:

- [Using Identifiers in a Safe Manner](#)

When you create a query through string concatenation, not only literals in queries pose a security risk, but also identifiers like graph names, labels, and property names do. The only problem is that bind variables are not supported for such identifier. Therefore, if these identifiers are variable from the application's perspective, then it is recommended to protect against query injection by passing the identifier through the

```
oracle.pgql.lang.ir.PgqlUtils.printIdentifier(String identifier)
method.
```

### 6.8.8.1 Using Bind Variables

There are two reasons for using bind variables:

- It protects against query injection.
- It speeds up query execution because the same bind variables can be set multiple times without requiring recompilation of the query.

To create a prepared statement, use one of the following two methods:

- `PgxGraph.preparePgql(String query) : PgxPreparedStatement`

- `PgxSession.preparePgql(String query) : PgxPreparedStatement`

The `PgxPreparedStatement` (package `oracle.pgx.api`) returned from these methods have setter methods for binding the bind variables to values of the designated data type.

```
PreparedStatement stmt = g.preparePgql(
 "SELECT v.id, v.dob " +
 "FROM MATCH (v) " +
 "WHERE v.firstName = ? AND v.lastName = ?");
stmt.setString(1, "Camille");
stmt.setString(2, "Mullins");
ResultSet rs = stmt.executeQuery();
```

Each bind variable in the query needs to be set to a value using one of the following setters of `PgxPreparedStatement`:

- `setBoolean(int parameterIndex, boolean x)`
- `setDouble(int parameterIndex, double x)`
- `setFloat(int parameterIndex, float x)`
- `setInt(int parameterIndex, int x)`
- `setLong(int parameterIndex, long x)`
- `setDate(int parameterIndex, LocalDate x)`
- `setTime(int parameterIndex, LocalTime x)`
- `setTimestamp(int parameterIndex, LocalDateTime x)`
- `setTimeWithTimezone(int parameterIndex, OffsetTime x)`
- `setTimestampWithTimezone(int parameterIndex, OffsetDateTime x)`
- `setArray(int parameterIndex, List<?> x)`

Once all the bind variables are set, the statement can be executed through:

- `PgxPreparedStatement.executeQuery()`
  - For `SELECT` queries only
  - Returns a `ResultSet`
- `PgxPreparedStatement.execute()`
  - For any type of statement
  - Returns a `Boolean` to indicate the form of the result: `true` in case of a `SELECT` query, `false` otherwise
  - In case of `SELECT`, the `ResultSet` can afterwards be accessed through `PgxPreparedStatement.getResultSet()`

In PGQL, bind variables can be used in place of literals of any data type, including array literals. An example query with a bind variable to is set to an instance of a `String` array is:

```
List<String> countryNames = new ArrayList<String>();
countryNames.add("Scotland");
countryNames.add("Tanzania");
```

```
countryNames.add("Serbia");

PreparedStatement stmt = g.preparePgql(
 "SELECT n.name, n.population " +
 "FROM MATCH (c:Country) " +
 "WHERE c.name IN ?");

ResultSet rs = stmt.executeQuery();
```

Finally, if a prepared statement is no longer needed, it is closed through `PgxPreparedStatement.close()` to free up resources.

### 6.8.8.2 Using Identifiers in a Safe Manner

When you create a query through string concatenation, not only literals in queries pose a security risk, but also identifiers like graph names, labels, and property names do. The only problem is that bind variables are not supported for such identifier. Therefore, if these identifiers are variable from the application's perspective, then it is recommended to protect against query injection by passing the identifier through the `oracle.pgql.lang.ir.PgqlUtils.printIdentifier(String identifier)` method.

Given an identifier string, the method automatically adds double quotes to the start and end of the identifier and escapes the characters in the identifier appropriately.

Consider the following example:

```
String graphNamePrinted = printIdentifier("my graph name with \"
special % characters ");
PreparedStatement stmt = g.preparePgql(
 "SELECT COUNT(*) AS numVertices FROM MATCH (v) ON " +
 graphNamePrinted);
```

### 6.8.9 Best Practices for Tuning PGQL Queries

This section describes best practices regarding memory allocation, parallelism, and query planning.

- **Memory Allocation**  
The In-Memory Analyst (PGX) has `on-heap` and `off-heap` memory, the earlier being the standard JVM heap while the latter being a separate heap that is managed by PGX. Just like graph data, intermediate and final results of PGQL queries are partially stored on-heap and partially off-heap. Therefore, both heaps are needed.
- **Parallelism**  
By default, all available processor threads are used to process PGQL queries. However, if needed, the number of threads can be limited by setting the `parallelism` option of the In-Memory Analyst (PGX).
- **Query Plan Explaining**  
The `PgxGraph.explainPgql(String query)` method is used to get insight into the query plan of the query. The method returns an instance of `Operation` (package `oracle.pgql.api`) which has the following methods:

### 6.8.9.1 Memory Allocation

The In-Memory Analyst (PGX) has `on-heap` and `off-heap` memory, the earlier being the standard JVM heap while the latter being a separate heap that is managed by PGX. Just like graph data, intermediate and final results of PGQL queries are partially stored on-heap and partially off-heap. Therefore, both heaps are needed.

In case of the on-heap memory, the default maximum is chosen upon startup of the JVM, but it can be overwritten through the `-Xmx` option.

In case of the off-heap, there is no maximum set by default and the off-heap memory usage, therefore, keeps increasing automatically until it depletes the system resources, in which case the operation is cancelled, its memory is released, and an appropriate exception is passed to the user. If needed, a maximum off-heap size can be configured through the `max_off_heap_size` option of PGX.

A ratio of 1:1 for on-heap vs. off-heap is recommended as a good starting point to allow for the largest possible graphs to be loaded and queried. For example, if you have 256 GB of memory available on your machine, then setting the maximum on-heap to 125 GB will make sure that there is a similar amount of memory available for off-heap:

```
export JAVA_OPTS="-Xmx125g"
```

### 6.8.9.2 Parallelism

By default, all available processor threads are used to process PGQL queries. However, if needed, the number of threads can be limited by setting the `parallelism` option of the In-Memory Analyst (PGX).

See [Configuration Parameters for the Graph Server \(PGX\) Engine](#) for more information on the graph server configuration parameters.

### 6.8.9.3 Query Plan Explaining

The `PgxGraph.explainPgql(String query)` method is used to get insight into the query plan of the query. The method returns an instance of `Operation` (package `oracle.pgx.api`) which has the following methods:

- `print()`: for printing the operation and its child operations
- `getOperationType()`: for getting the type of the operation
- `getPatternInfo()`: for getting a string representation of the operation
- `getCostEstimate()`: for getting the cost of the operation
- `getTotalCostEstimate()`: for getting the cost of the operations and its child operations
- `getCardinalityEstimate()`: for getting the expected number of result rows
- `getChildren()`: for accessing the child operations

Consider the following example:

```
g.explainPgql("SELECT COUNT(*) FROM MATCH (n) -[e1]-> (m) -[e2]-> (o)").print()
\--- GROUP BY GroupBy {"cardinality":"42", "cost":"42",
"accumulatedCost":"58.1"}
```

```

 \--- (m) -[e2]-> (o) NeighborMatch {"cardinality":"3.12",
"cost":"3.12", "accumulatedCost":"16.1"}
 \--- (n) -[e1]-> (m) NeighborMatch {"cardinality":"5",
"cost":"5", "accumulatedCost":"13"}
 \--- (n) RootVertexMatch {"cardinality":"8",
"cost":"8", "accumulatedCost":"8"}

```

In the above example, the `print()` method is used to print the query plan.

If a query plan is not optimal, it is often possible to rewrite the query to improve its performance. For example, a `SELECT` query may be split into an `UPDATE` and a `SELECT` query as a way to improve the total runtime.

Note that the In-Memory Analyst (PGX) does not provide a hint mechanism.

## 6.9 Executing PGQL Queries Directly Against Oracle Database

This topic explains how you can execute PGQL queries directly against the graph in Oracle Database (as opposed to in-memory).

Property Graph Query Language (PGQL) queries can be executed against disk-resident property graph data stored in Oracle Database. PGQL on Oracle Database (RDBMS) provides a Java API for executing PGQL queries. Logic in PGQL on RDBMS translates a submitted PGQL query into an equivalent SQL query, and the resulting SQL is executed on the database server. PGQL on RDBMS then wraps the SQL query results with a convenient PGQL result set API.

Property graph data in RDBMS can exist either in the property graph schema tables or as a property graph view on the Oracle Database tables.

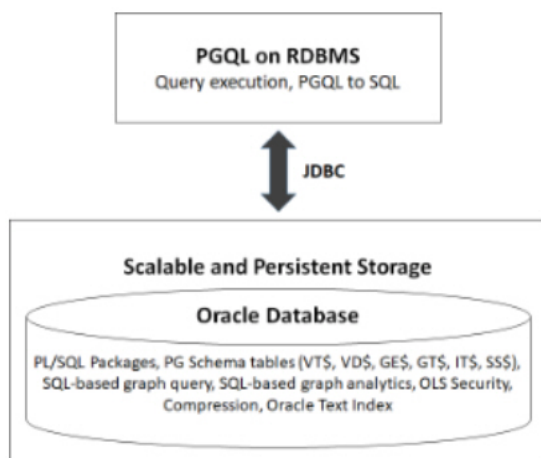
The following topics explain in detail how you can execute PGQL queries against the graph in the Oracle Database:

- [Executing PGQL Queries Against Property Graph Schema Tables](#)  
This topic explains how you can execute PGQL queries directly against the graph stored in property graph schema tables.
- [Executing PGQL Queries Against Property Graph Views](#)  
This topic explains how you can execute PGQL queries directly against the property graph views on Oracle Database tables.

### 6.9.1 Executing PGQL Queries Against Property Graph Schema Tables

This topic explains how you can execute PGQL queries directly against the graph stored in property graph schema tables.

The PGQL query execution flow is shown in the following figure.

**Figure 6-1 PGQL on Property Graph Schema Tables in Oracle Database (RDBMS)**

The basic execution flow is:

1. The PGQL query is submitted to PGQL on RDBMS through a Java API.
2. The PGQL query is translated to SQL.
3. The translated SQL is submitted to Oracle Database by JDBC.
4. The SQL result set is wrapped as a PGQL result set and returned to the caller.

The ability to execute PGQL queries directly against property graph data stored in Oracle Database provides several benefits.

- PGQL provides a more natural way to express graph queries than SQL manually written to query schema tables, including VT\$, VD\$, GE\$, and GT\$.
- PGQL queries can be executed without the need to load a snapshot of your graph data into PGX, so there is no need to worry about staleness of frequently updated graph data.
- PGQL queries can be executed against graph data that is too large to fit in memory.
- The robust and scalable Oracle SQL engine can be used to execute PGQL queries.
- Mature tools for management, monitoring and tuning of Oracle Database can be used to tune and monitor PGQL queries.
- [PGQL Features Supported in Property Graph Schema](#)
- [Creating Property Graphs through CREATE PROPERTY GRAPH Statements](#)
- [Dropping Property Graphs through DROP PROPERTY GRAPH Statements](#)
- [Using the oracle.pg.rdbms.pgql Java Package to Execute PGQL Queries](#)
- [Using the Python Client to Execute PGQL Queries](#)
- [Performance Considerations for PGQL Queries](#)

### 6.9.1.1 PGQL Features Supported in Property Graph Schema

**PGQL** is a SQL-like query language for querying property graph data. It is based on the concept of graph pattern matching and allows you to specify, among other things, topology constraints, paths, filters, sorting and aggregation.

The Java API for PGQL defined in the `oracle.pg.rdbms.pgql` package supports the PGQL specification with a few exceptions. (The PGQL specification can be found here - [PGQL Specification](#)).

The following are a few PGQL features which are not supported:

- Shortest path
- ARRAY\_AGG aggregation
- Single CHEAPEST path and TOP-K CHEAPEST path using `COST` functions
- Case-insensitive matching of uppercased references to labels and properties



#### Note:

See [Supported PGQL Features and Limitations](#) for a complete list of supported and unsupported PGQL features for PGQL on RDBMS (PG\_SCHEMA option) .

The following explains a few PGQL features that require special consideration.

- [Temporal Types](#)
- [Type Casting](#)
- [CONTAINS Built-in Function](#)

### 6.9.1.1.1 Temporal Types

The temporal types DATE, TIMESTAMP and TIMESTAMP WITH TIMEZONE are supported in PGQL queries.

All of these value types are represented internally using the Oracle SQL TIMESTAMP WITH TIME ZONE type. DATE values are automatically converted to TIMESTAMP WITH TIME ZONE by assuming the earliest time in UTC+0 timezone (for example, 2000-01-01 becomes 2000-01-01 00:00:00.00+00:00). TIMESTAMP values are automatically converted to TIMESTAMP WITH TIME ZONE by assuming UTC+0 timezone (for example, 2000-01-01 12:00:00.00 becomes 2000-01-01 12:00:00.00+00:00).

Temporal constants are written in PGQL queries as follows.

- DATE 'YYYY-MM-DD'
- TIMESTAMP 'YYYY-MM-DD HH24:MI:SS.FF'
- TIMESTAMP WITH TIMEZONE 'YYYY-MM-DD HH24:MI:SS.FFTZH:TZM'

Some examples are DATE '2000-01-01', TIMESTAMP '2000-01-01 14:01:45.23', TIMESTAMP WITH TIMEZONE '2000-01-01 13:00:00.00-05:00', and TIMESTAMP WITH TIMEZONE '2000-01-01 13:00:00.00+01:00'.

In addition, temporal values can be obtained by casting string values to a temporal type. The supported string formats are:

- DATE 'YYYY-MM-DD'
- TIMESTAMP 'YYYY-MM-DD HH24:MI:SS.FF' and 'YYYY-MM-DD"T"HH24:MI:SS.FF'



- `TIMESTAMP WITH TIMEZONE 'YYYY-MM-DD HH24:MI:SS.FFTZH:TZM'` and `'YYYY-MM-DD"T"HH24:MI:SS.FFTZH:TZM'`.

Some examples are `CAST ('2005-02-04' AS DATE)`, `CAST ('1990-01-01 12:00:00.00' AS TIMESTAMP)`, `CAST ('1985-01-01T14:05:05.00-08:00' AS TIMESTAMP WITH TIMEZONE)`.

When consuming results from a `PgqlResultSet` object, `getObject` returns a `java.sql.Timestamp` object for temporal types.

Bind variables can only be used for the `TIMESTAMP WITH TIMEZONE` temporal type in PGQL, and a `setTimestamp` method that takes a `java.sql.Timestamp` object as input is used to set the bind value. As a simpler alternative, you can use a string bind variable in a `CAST` statement to bind temporal values (for example, `CAST (? AS TIMESTAMP WITH TIMEZONE)` followed by `setString(1, "1985-01-01T14:05:05.00-08:00")`). See also [Using Bind Variables in PGQL Queries](#) for more information about bind variables.

### 6.9.1.1.2 Type Casting

Type casting is supported in PGQL with a SQL-style `CAST (VALUE AS DATATYPE)` syntax, for example `CAST('25' AS INT)`, `CAST (10 AS STRING)`, `CAST ('2005-02-04' AS DATE)`, `CAST(e.weight AS STRING)`. Supported casting operations are summarized in the following table. Y indicates that the conversion is supported, and N indicates that it is not supported. Casting operations on invalid values (for example, `CAST('xyz' AS INT)`) or unsupported conversions (for example, `CAST (10 AS TIMESTAMP)`) return NULL instead of raising a SQL exception.

**Table 6-3 Type Casting Support in PGQL (From and To Types)**

| "to" type                            | from<br>STRIN<br>G | from<br>INT | from<br>LONG | from<br>FLOA<br>T | from<br>DOUBL<br>E | from<br>BOOLEA<br>N | from<br>DATE | from<br>TIMESTAM<br>P | from<br>TIMESTAM<br>P WITH<br>TIMEZONE |
|--------------------------------------|--------------------|-------------|--------------|-------------------|--------------------|---------------------|--------------|-----------------------|----------------------------------------|
| to STRING                            | Y                  | Y           | Y            | Y                 | Y                  | Y                   | Y            | Y                     | Y                                      |
| to INT                               | Y                  | Y           | Y            | Y                 | Y                  | Y                   | N            | N                     | N                                      |
| to LONG                              | Y                  | Y           | Y            | Y                 | Y                  | Y                   | N            | N                     | N                                      |
| to FLOAT                             | Y                  | Y           | Y            | Y                 | Y                  | Y                   | N            | N                     | N                                      |
| to DOUBLE                            | Y                  | Y           | Y            | Y                 | Y                  | Y                   | N            | N                     | N                                      |
| to<br>BOOLEAN                        | Y                  | Y           | Y            | Y                 | Y                  | Y                   | N            | N                     | N                                      |
| to DATE                              | Y                  | N           | N            | N                 | N                  | N                   | Y            | Y                     | Y                                      |
| to<br>TIMESTAM<br>P                  | Y                  | N           | N            | N                 | N                  | N                   | Y            | Y                     | Y                                      |
| to<br>TIMESTAM<br>P WITH<br>TIMEZONE | Y                  | N           | N            | N                 | N                  | N                   | Y            | Y                     | Y                                      |

An example query that uses type casting is:

```
SELECT e.name, CAST (e.birthDate AS STRING) AS dob
FROM MATCH (e)
WHERE e.birthDate < CAST ('1980-01-01' AS DATE)
```

### 6.9.1.1.3 CONTAINS Built-in Function

A CONTAINS built-in function is supported. It is used in conjunction with an Oracle Text index on vertex and edge properties. CONTAINS returns `true` if a value matches an Oracle Text search string and `false` if it does not match.

An example query is:

```
SELECT v.name
FROM MATCH (v)
WHERE CONTAINS(v.abstract, 'Oracle')
```

See also [Using a Text Index with PGQL Queries](#) for more information about using full text indexes with PGQL.

## 6.9.1.2 Creating Property Graphs through CREATE PROPERTY GRAPH Statements

You can use PGQL to create property graphs from relational database tables. A CREATE PROPERTY GRAPH statement defines a set of vertex tables that are transformed into vertices and a set of edge tables that are transformed into edges. For each table a key, a label and a set of column properties can be specified. The column types CHAR, NCHAR, VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, LONG, FLOAT, DATE, TIMESTAMP and TIMESTAMP WITH TIMEZONE are supported for CREATE PROPERTY GRAPH column properties.

When a CREATE PROPERTY GRAPH statement is called, a property graph schema for the graph is created, and the data is copied from the source tables into the property graph schema tables. The graph is created as a one-time copy and is not automatically kept in sync with the source data.

### Example 6-3 PgqlCreateExample1.java

This example shows how to create a property graph from a set of relational tables. Notice that the example creates tables Person, Hobby, and Hobbies, so they should not exist before running the example. The example also shows how to execute a query against a property graph.

```
import java.sql.Connection;
import java.sql.Statement;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;
```

```

/**
 * This example shows how to create a Property Graph from relational
 * data stored in Oracle Database executing a PGQL create statement.
 */
public class PgqlCreateExample1
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 Statement stmt = null;
 PgqlStatement pgqlStmt = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port +":"+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Create relational data
 stmt = conn.createStatement();

 //Table Person
 stmt.executeUpdate(
 "create table Person(" +
 " id NUMBER, " +
 " name VARCHAR2(20), " +
 " dob TIMESTAMP " +
 ")");

 // Insert some data
 stmt.executeUpdate("insert into Person values(1,'Alan', DATE
'1995-05-26')");
 stmt.executeUpdate("insert into Person values(2,'Ben', DATE
'2007-02-15')");
 stmt.executeUpdate("insert into Person values(3,'Claire', DATE
'1967-11-30')");

 // Table Hobby
 stmt.executeUpdate(
 "create table Hobby(" +

```

```

" id NUMBER, " +
" name VARCHAR2(20) " +
"");

// Insert some data
stmt.executeUpdate("insert into Hobby values(1, 'Sports')");
stmt.executeUpdate("insert into Hobby values(2, 'Music')");

// Table Hobbies
stmt.executeUpdate(
"create table Hobbies("+
" person NUMBER, "+
" hobby NUMBER, "+
" strength NUMBER "+
"");

// Insert some data
stmt.executeUpdate("insert into Hobbies values(1, 1, 20)");
stmt.executeUpdate("insert into Hobbies values(1, 2, 30)");
stmt.executeUpdate("insert into Hobbies values(2, 1, 10)");
stmt.executeUpdate("insert into Hobbies values(3, 2, 20)");

//Commit changes
conn.commit();

// Get a PGQL connection
PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);

// Create a PgqlStatement
pgqlStmt = pgqlConn.createStatement();

// Execute PGQL to create property graph
String pgql =
"Create Property Graph " + graph + " " +
"VERTEX TABLES (" +
" Person " +
" Key(id) " +
" Label \"people\"" +
" PROPERTIES(name AS \"first_name\", dob AS \"birthday\"),"
+
" Hobby " +
" Key(id) Label \"hobby\" PROPERTIES(name AS \"name\"" +
")" +
"EDGE TABLES (" +
" Hobbies" +
" SOURCE KEY(person) REFERENCES Person " +
" DESTINATION KEY(hobby) REFERENCES Hobby " +
" LABEL \"likes\" PROPERTIES (strength AS \"score\"" +
");
pgqlStmt.execute(pgql);

// Execute a PGQL query to verify Graph creation
pgql =
"SELECT p.\"first_name\", p.\"birthday\", h.\"name\",
e.\"score\" " +

```

```

 "FROM MATCH (p:\"people\")-[e:\"likes\"]->(h:\"hobby\") ON " + graph;
rs = pgqlStmt.executeQuery(pgql, "");

 // Print the results
 rs.print();
}
finally {
 // close the sql statement
 if (stmt != null) {
 stmt.close();
 }
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (pgqlStmt != null) {
 pgqlStmt.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
}
}
}
}
}

```

The output for `PgqlCreateExample1.java` is:

```

+-----+
| first_name | birthday | name | score |
+-----+
Alan	1995-05-25 17:00:00.0	Music	30.0
Claire	1967-11-29 16:00:00.0	Music	20.0
Ben	2007-02-14 16:00:00.0	Sports	10.0
Alan	1995-05-25 17:00:00.0	Sports	20.0
+-----+

```

#### Example 6-4 PgqlCreateExample2.java

This example shows how a create property graph statement without specifying any keys. Notice that the example creates tables `Person`, `Hobby`, and `Hobbies`, so they should not exist before running the example.

```

import java.sql.Connection;
import java.sql.Statement;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to create a Property Graph from relational

```

```

* data stored in Oracle Database executing a PGQL create statement.
*/
public class PgqlCreateExample2
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 Statement stmt = null;
 PgqlStatement pgqlStmt = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port +":"+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Create relational data
 stmt = conn.createStatement();

 //Table Person
 stmt.executeUpdate(
 "create table Person(" +
 " id NUMBER, " +
 " name VARCHAR2(20), " +
 " dob TIMESTAMP, " +
 " CONSTRAINT pk_person PRIMARY KEY(id) " +
 ")");

 // Insert some data
 stmt.executeUpdate("insert into Person values(1,'Alan', DATE
'1995-05-26')");
 stmt.executeUpdate("insert into Person values(2,'Ben', DATE
'2007-02-15')");
 stmt.executeUpdate("insert into Person values(3,'Claire', DATE
'1967-11-30')");

 // Table Hobby
 stmt.executeUpdate(
 "create table Hobby(" +

```

```

" id NUMBER, " +
" name VARCHAR2(20), " +
" CONSTRAINT pk_hobby PRIMARY KEY(id) " +
"");

// Insert some data
stmt.executeUpdate("insert into Hobby values(1, 'Sports')");
stmt.executeUpdate("insert into Hobby values(2, 'Music')");

// Table Hobbies
stmt.executeUpdate(
"create table Hobbies("+
" person NUMBER, "+
" hobby NUMBER, "+
" strength NUMBER, "+
" CONSTRAINT fk_hobbies1 FOREIGN KEY (person) REFERENCES
Person(id), "+
" CONSTRAINT fk_hobbies2 FOREIGN KEY (hobby) REFERENCES Hobby(id)"+
"");

// Insert some data
stmt.executeUpdate("insert into Hobbies values(1, 1, 20)");
stmt.executeUpdate("insert into Hobbies values(1, 2, 30)");
stmt.executeUpdate("insert into Hobbies values(2, 1, 10)");
stmt.executeUpdate("insert into Hobbies values(3, 2, 20)");

//Commit changes
conn.commit();

// Get a PGQL connection
PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);

// Create a PgqlStatement
pgqlStmt = pgqlConn.createStatement();

// Execute PGQL to create property graph
String pgql =
"Create Property Graph " + graph + " " +
"VERTEX TABLES (" +
" Person " +
" Label people +
" PROPERTIES ALL COLUMNS," +
" Hobby " +
" Label hobby PROPERTIES ALL COLUMNS EXCEPT(id) " +
")" +
"EDGE TABLES (" +
" Hobbies" +
" SOURCE Person DESTINATION Hobby " +
" LABEL likes NO PROPERTIES" +
");";
pgqlStmt.executeUpdate(pgql);

// Execute a PGQL query to verify Graph creation
pgql =
"SELECT p.NAME AS person, p.DOB, h.NAME AS hobby " +

```

```

 "FROM MATCH (p:people)-[e:likes]->(h:hobby) ON " + graph;
rs = pgqlStmt.executeQuery(pgql, "");

// Print the results
rs.print();
}
finally {
// close the sql statement
if (stmt != null) {
 stmt.close();
}
// close the result set
if (rs != null) {
 rs.close();
}
// close the statement
if (pgqlStmt != null) {
 pgqlStmt.close();
}
// close the connection
if (conn != null) {
 conn.close();
}
}
}
}
}

```

The output for `PgqlCreateExample2.java` is:

```

+-----+
| PERSON | DOB | HOBBY |
+-----+
Alan	1995-05-25 17:00:00.0	Music
Claire	1967-11-29 16:00:00.0	Music
Ben	2007-02-14 16:00:00.0	Sports
Alan	1995-05-25 17:00:00.0	Sports
+-----+

```

### 6.9.1.3 Dropping Property Graphs through DROP PROPERTY GRAPH Statements

You can use PGQL to drop property graphs. When a `DROP PROPERTY GRAPH` statement is called, all the property graph schema tables of the graph are dropped.

#### Example 6-5 `PgqlDropExample1.java`

This example shows how to drop a property graph.

```

import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

```



```
/**
 * This example shows how to drop a Property executing a PGQL drop statement.
 */
public class PgqlDropExample1
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement pgqlStmt = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@" + host + ":" + port + ":" + sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);

 // Create a PgqlStatement
 pgqlStmt = pgqlConn.createStatement();

 // Execute PGQL to drop property graph
 String pgql = "Drop Property Graph " + graph;
 pgqlStmt.execute(pgql);

 }
 finally {
 // close the statement
 if (pgqlStmt != null) {
 pgqlStmt.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
 }
 }
}
```

### 6.9.1.4 Using the `oracle.pg.rdbms.pgql` Java Package to Execute PGQL Queries

The Java API in the `oracle.pg.rdbms.pgql` package provides support for executing PGQL queries against Oracle Database. This topic explains how to use the Java API through a series of examples.

 **Note:**

Effective with Release 21c, the following classes in the `oracle.pg.rdbms` package are deprecated:

```
oracle.pg.rdbms.OraclePgqlColumnDescriptorImpl
oracle.pg.rdbms.OraclePgqlColumnDescriptor
oracle.pg.rdbms.OraclePgqlExecutionFactory
oracle.pg.rdbms.OraclePgqlExecution
oracle.pg.rdbms.PgqlPreparedStatement
oracle.pg.rdbms.OraclePgqlResultElementImpl
oracle.pg.rdbms.OraclePgqlResultElement
oracle.pg.rdbms.OraclePgqlResultImpl
oracle.pg.rdbms.OraclePgqlResultIterable
oracle.pg.rdbms.OraclePgqlResultIteratorImpl
oracle.pg.rdbms.OraclePgqlResult
oracle.pg.rdbms.OraclePgqlResultSetImpl
oracle.pg.rdbms.OraclePgqlResultSet
oracle.pg.rdbms.OraclePgqlResultSetMetaDataImpl
oracle.pg.rdbms.OraclePgqlResultSetMetaData
oracle.pg.rdbms.PgqlSqlQueryTransImpl
oracle.pg.rdbms.PgqlSqlQueryTrans
oracle.pg.rdbms.PgqlStatement
```

You should instead use equivalent classes in `oracle.pg.rdbms.pgql`:

```
oracle.pg.rdbms.pgql.PgqlColumnDescriptorImpl
oracle.pg.rdbms.pgql.PgqlColumnDescriptor
oracle.pg.rdbms.pgql.PgqlConnection
oracle.pg.rdbms.pgql.PgqlExecution
oracle.pg.rdbms.pgql.PgqlPreparedStatement
oracle.pg.rdbms.pgql.PgqlResultElementImpl
oracle.pg.rdbms.pgql.PgqlResultElement
oracle.pg.rdbms.pgql.PgqlResultSetImpl
oracle.pg.rdbms.pgql.PgqlResultSet
oracle.pg.rdbms.pgql.PgqlResultSetMetaDataImpl
oracle.pg.rdbms.pgql.PgqlSqlTransImpl
oracle.pg.rdbms.pgql.PgqlSqlTrans
oracle.pg.rdbms.pgql.PgqlStatement
```

One difference between `oracle.pg.rdbms.OraclePgqlResultSet` and `oracle.pg.rdbms.pgql.PgqlResultSet` is that `oracle.pg.rdbms.pgql.PgqlResultSet` does not provide APIs to retrieve vertex and edge objects. Existing code using those interfaces should be changed to project IDs rather than `OracleVertex` and `OracleEdge` objects. You can obtain an `OracleVertex` or `OracleEdge` object from the projected ID values by calling `OracleVertex.getInstance()` or `OracleEdge.getInstance()`. (For an example, see [Example 6-21](#).)

See [Oracle Graph Property Graph Java APIs](#) for more details on setting the classpath for compiling and executing your Java applications.

The following `test_graph` data set in Oracle flat file format will be used in the examples in subtopics that follow. The data set includes a vertex file (`test_graph.opv`) and an edge file (`test_graph.ope`).

`test_graph.opv`:

```
2, fname, 1, Ray, , , person
2, lname, 1, Green, , , person
2, mval, 5, , , 1985-01-01T12:00:00.000Z, person
2, age, 2, , 41, , person
0, bval, 6, Y, , , person
0, fname, 1, Bill, , , person
0, lname, 1, Brown, , , person
0, mval, 1, y, , , person
0, age, 2, , 40, , person
1, bval, 6, Y, , , person
1, fname, 1, John, , , person
1, lname, 1, Black, , , person
1, mval, 2, , 27, , person
1, age, 2, , 30, , person
3, bval, 6, N, , , person
3, fname, 1, Susan, , , person
3, lname, 1, Blue, , , person
3, mval, 6, N, , , person
3, age, 2, , 35, , person
```

test\_graph.ope:

```
4, 0, 1, knows, mval, 1, Y, ,
4, 0, 1, knows, firstMetIn, 1, MI, ,
4, 0, 1, knows, since, 5, , , 1990-01-01T12:00:00.000Z
16, 0, 1, friendOf, strength, 2, , 6,
7, 1, 0, knows, mval, 5, , , 2003-01-01T12:00:00.000Z
7, 1, 0, knows, firstMetIn, 1, GA, ,
7, 1, 0, knows, since, 5, , , 2000-01-01T12:00:00.000Z
17, 1, 0, friendOf, strength, 2, , 7,
9, 1, 3, knows, mval, 6, N, ,
9, 1, 3, knows, firstMetIn, 1, SC, ,
9, 1, 3, knows, since, 5, , , 2005-01-01T12:00:00.000Z
10, 2, 0, knows, mval, 1, N, ,
10, 2, 0, knows, firstMetIn, 1, TX, ,
10, 2, 0, knows, since, 5, , , 1997-01-01T12:00:00.000Z
12, 2, 3, knows, mval, 3, , 342.5,
12, 2, 3, knows, firstMetIn, 1, TX, ,
12, 2, 3, knows, since, 5, , , 2011-01-01T12:00:00.000Z
19, 2, 3, friendOf, strength, 2, , 4,
14, 3, 1, knows, mval, 1, a, ,
14, 3, 1, knows, firstMetIn, 1, CA, ,
14, 3, 1, knows, since, 5, , , 2010-01-01T12:00:00.000Z
15, 3, 2, knows, mval, 1, z, ,
15, 3, 2, knows, firstMetIn, 1, CA, ,
15, 3, 2, knows, since, 5, , , 2004-01-01T12:00:00.000Z
5, 0, 2, knows, mval, 2, , 23,
5, 0, 2, knows, firstMetIn, 1, OH, ,
5, 0, 2, knows, since, 5, , , 2002-01-01T12:00:00.000Z
6, 0, 3, knows, mval, 3, , 159.7,
6, 0, 3, knows, firstMetIn, 1, IN, ,
6, 0, 3, knows, since, 5, , , 1994-01-01T12:00:00.000Z
8, 1, 2, knows, mval, 6, Y, ,
8, 1, 2, knows, firstMetIn, 1, FL, ,
8, 1, 2, knows, since, 5, , , 1999-01-01T12:00:00.000Z
18, 1, 3, friendOf, strength, 2, , 5,
11, 2, 1, knows, mval, 2, , 1001,
11, 2, 1, knows, firstMetIn, 1, OK, ,
11, 2, 1, knows, since, 5, , , 2003-01-01T12:00:00.000Z
13, 3, 0, knows, mval, 5, , , 2001-01-01T12:00:00.000Z
13, 3, 0, knows, firstMetIn, 1, CA, ,
```

```
13,3,0, knows, since, 5,,, 2006-01-01T12:00:00.000Z
20,3,1, friendOf, strength, 2,, 3,
```

- [Basic Query Execution](#)
- [Executing PGQL Queries Using JDBC Driver](#)
- [Security Techniques for PGQL Queries](#)
- [Using a Text Index with PGQL Queries](#)
- [Obtaining the SQL Translation for a PGQL Query](#)
- [Additional Options for PGQL Translation and Execution](#)
- [Querying Another User's Property Graph](#)
- [Using Query Optimizer Hints with PGQL](#)
- [Modifying Property Graphs through INSERT, UPDATE, and DELETE Statements](#)

#### 6.9.1.4.1 Basic Query Execution

Two main Java Interfaces, `PgqlStatement` and `PgqlResultSet`, are used for PGQL execution. This topic includes several examples of basic query execution.

##### Example 6-6 `GraphLoaderExample.java`

`GraphLoaderExample.java` loads some Oracle property graph data that will be used in subsequent examples in this topic.

```
import oracle.pg.rdbms.Oracle;
import oracle.pg.rdbms.OraclePropertyGraph;
import oracle.pg.rdbms.OraclePropertyGraphDataLoader;

/**
 * This example shows how to create an Oracle Property Graph
 * and load data into it from vertex and edge flat files.
 */
public class GraphLoaderExample
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];
 String vertexFile = args[idx++];
 String edgeFile = args[idx++];

 Oracle oracle = null;
 OraclePropertyGraph opg = null;

 try {
 // Create a connection to Oracle
 oracle = new Oracle("jdbc:oracle:thin:@" + host + ":" + port + ":" + sid, user,
```

```
password);

 // Create a property graph
 opg = OraclePropertyGraph.getInstance(oracle, graph);

 // Clear any existing data
 opg.clearRepository();

 // Load data from opv and ope files
 OraclePropertyGraphDataLoader opgLoader =
OraclePropertyGraphDataLoader.getInstance();
 opgLoader.loadData(opg, vertexFile, edgeFile, 1);

 System.out.println("Vertices loaded:" + opg.countVertices());
 System.out.println("Edges loaded:" + opg.countEdges());

}
finally {
 // close the property graph
 if (opg != null) {
 opg.close();
 }
 // close oracle
 if (oracle != null) {
 oracle.dispose();
 }
}
}
}
```

GraphLoaderExample.java gives the following output for test\_graph.

```
Vertices loaded:4
Edges loaded:17
```

### Example 6-7 PgqlExample1.java

PgqlExample1.java executes a PGQL query and prints the query result. PgqlConnection is used to obtain a PgqlStatement. Next, it calls the executeQuery method of PgqlStatement, which returns a PgqlResultSet object. PgqlResultSet provides a print() method, which shows results in a tabular mode.

The PgqlResultSet and PgqlStatement objects should be closed after consuming the query result.

```
import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
```

```

* This example shows how to execute a basic PGQL query against disk-
resident
* PG data stored in Oracle Database and iterate through the result.
*/
public class PgqlExample1
{

public static void main(String[] args) throws Exception
{
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port +" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Execute query to get a PgqlResultSet object
 String pgql =
mval "+
 "SELECT v.\"fname\" AS fname, v.\"lname\" AS lname, v.\"mval\" AS
 "FROM MATCH (v)";
 rs = ps.executeQuery(pgql, /* query string */
 "" /* options */);

 // Print the results
 rs.print();
 }
 finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 }
}
}

```

```

 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
 }
}
}
}

```

PgqlExample1.java gives the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```

+-----+
| FNAME | LNAME | MVAL |
+-----+
Susan	Blue	false
Bill	Brown	y
Ray	Green	1985-01-01 04:00:00.0
John	Black	27
+-----+

```

### Example 6-8 PgqlExample2.java

PgqlExample2.java shows a PGQL query with a temporal filter on an edge property.

- PgqlResultSet provides an interface for consuming the query result that is very similar to the java.sql.ResultSet interface.
- A next() method allows moving through the query result, and a close() method allows releasing resources after the application is finished reading the query result.
- In addition, PgqlResultSet provides getters for String, Integer, Long, Float, Double, Boolean, LocalDateTime, and OffsetDateTime, and it provides a generic getObject() method for values of any type.

```

import java.sql.Connection;

import java.text.SimpleDateFormat;

import java.util.Date;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.pgql.lang.ResultSet;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to execute a PGQL query with a temporal edge
 * property filter against disk-resident PG data stored in Oracle
 * Database
 * and iterate through the result.
 */

```



```

public class PgqlExample2
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 ResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port +":"+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Create a Pgql connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Execute query to get a ResultSet object
 String pgql =
loc "+
 "SELECT v.\"fname\" AS n1, v2.\"fname\" AS n2, e.\"firstMetIn\" AS
 "FROM MATCH (v)-[e:\"knows\"]->(v2) "+
 "WHERE e.\"since\" > TIMESTAMP '2000-01-01 00:00:00.00+00:00'";
 rs = ps.executeQuery(pgql, "");

 // Print results
 printResults(rs);
 }
 finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 }
 }
}

```

```

 if (conn != null) {
 conn.close();
 }
 }
}

/**
 * Prints a PGQL ResultSet
 */
static void printResults(ResultSet rs) throws Exception
{
 StringBuffer buff = new StringBuffer("");
 SimpleDateFormat sdf = new SimpleDateFormat("yyyy-MM-
dd'T'HH:mm:ss.SSSXXX");
 while (rs.next()) {
 buff.append("[");
 for (int i = 1; i <= rs.getMetaData().getColumnCount(); i++) {
 // use generic getObject to handle all types
 Object mval = rs.getObject(i);
 String mStr = "";
 if (mval instanceof java.lang.String) {
 mStr = "STRING: "+mval.toString();
 }
 else if (mval instanceof java.lang.Integer) {
 mStr = "INTEGER: "+mval.toString();
 }
 else if (mval instanceof java.lang.Long) {
 mStr = "LONG: "+mval.toString();
 }
 else if (mval instanceof java.lang.Float) {
 mStr = "FLOAT: "+mval.toString();
 }
 else if (mval instanceof java.lang.Double) {
 mStr = "DOUBLE: "+mval.toString();
 }
 else if (mval instanceof java.sql.Timestamp) {
 mStr = "DATE: "+sdf.format((Date)mval);
 }
 else if (mval instanceof java.lang.Boolean) {
 mStr = "BOOLEAN: "+mval.toString();
 }
 if (i > 1) {
 buff.append(",\t");
 }
 buff.append(mStr);
 }
 buff.append("]\n");
 }
 System.out.println(buff.toString());
}
}

```

PgqlExample2.java gives the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```
[STRING: Susan, STRING: Bill, STRING: CA]
[STRING: Susan, STRING: John, STRING: CA]
[STRING: Susan, STRING: Ray, STRING: CA]
[STRING: Bill, STRING: Ray, STRING: OH]
[STRING: Ray, STRING: John, STRING: OK]
[STRING: Ray, STRING: Susan, STRING: TX]
[STRING: John, STRING: Susan, STRING: SC]
[STRING: John, STRING: Bill, STRING: GA]
```

### Example 6-9 PgqlExample3.java

PgqlExample3.java shows a PGQL query with grouping and aggregation.

```
import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to execute a PGQL query with aggregation
 * against disk-resident PG data stored in Oracle Database and iterate
 * through the result.
 */
public class PgqlExample3
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {
 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@" + host + ":" + port + ":" + sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Create a Pgql connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);
```

```

// Create a PgqlStatement
ps = pgqlConn.createStatement();

// Execute query to get a ResultSet object
String pgql =
 "SELECT v.\"fname\" AS \"fname\", COUNT(v2) AS \"friendCnt\" "+
 "FROM MATCH (v)-[e:\"friendOf\"]->(v2) "+
 "GROUP BY v "+
 "ORDER BY \"friendCnt\" DESC";
rs = ps.executeQuery(pgql, "");

// Print results
rs.print();
}
finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
}
}

```

PgqlExample3.java gives the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```

+-----+
| fname | friendCnt |
+-----+
John	2
Bill	1
Ray	1
Susan	1
+-----+

```

#### Example 6-10 PgqlExample4.java

PgqlExample4.java shows a PGQL path query.

```

import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

```

```

/**
 * This example shows how to execute a path query in PGQL against
 * disk-resident PG data stored in Oracle Database and iterate
 * through the result.
 */
public class PgqlExample4
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+":"+port+"@"+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Create a Pgql connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Execute query to get a ResultSet object
 String pgql =
 "PATH fof AS ()-[:\"friendOf\"|\"knows\"]->() "+
 "SELECT v2.\"fname\" AS friend "+
 "FROM MATCH (v)-/:fof*/->(v2) "+
 "WHERE v.\"fname\" = 'John' AND v != v2";
 rs = ps.executeQuery(pgql, "");

 // Print results
 rs.print();
 }
 finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 }
 }
}

```

```

 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
}

```

PgqlExample4.java gives the following output for test\_graph(which can be loaded using GraphLoaderExample.java code).

```

+-----+
| FRIEND |
+-----+
| Susan |
| Bill |
| Ray |
+-----+

```

#### 6.9.1.4.2 Executing PGQL Queries Using JDBC Driver

The Oracle Graph Server and Client Release 21.2.0 includes a JDBC driver which allows you to run PGQL queries directly against the Oracle Database. To use the driver, register the following class at the JDBC driver manager:

```

import java.sql.DriverManager;
import oracle.pg.rdbms.pgql.jdbc.PgqlJdbcRdbmsDriver;
...
DriverManager.registerDriver(new PgqlJdbcRdbmsDriver());

```

To make JDBC use the driver, you need to prefix the JDBC URLs with jdbc:oracle:pgql as shown in this example:

```

import java.sql.Connection;
import java.sql.DriverManager;

Connection conn = DriverManager.getConnection("jdbc:oracle:pgql:@<DB
Host>:<DB Port>/<DB SID>", "<DB Username>", "<DB Password>");

```

The part after jdbc:oracle:pgql follows the same syntax as the regular Oracle JDBC thin driver. In other words, you can convert any valid Oracle JDBC thin driver URL into a PGQL driver URL by replacing jdbc:oracle:thin with jdbc:oracle:pgql. Once you obtained a connection object, you can use it to query property graphs using PGQL syntax. For example:

##### **Example 6-11 Executing a PGQL Query using the PGQL JDBC driver**

```

import java.sql.Connection;
import java.sql.DriverManager;

```

```

import java.sql.ResultSet;
import java.sql.PreparedStatement;
import oracle.pg.rdbms.pgql.jdbc.PgqlJdbcRdbmsDriver;

public class PgqlJdbcTest {

 public static void main(String[] args) throws Exception {

 DriverManager.registerDriver(new PgqlJdbcRdbmsDriver());
 String jdbcUrl = "jdbc:oracle:pgql:@<DB Host>:<DB Port>/<DB SID>";
 String username = "<DB Username>";
 String password = "<DB Password>";

 try (Connection conn = DriverManager.getConnection(jdbcUrl, username,
password)) {
 String query = "SELECT n.name FROM MATCH(n) ON test_graph WHERE id(n)
= ?";
 PreparedStatement pstmt = conn.prepareStatement(query);
 pstmt.setLong(1, 10L);
 pstmt.execute();
 ResultSet rs = pstmt.getResultSet();
 while(rs.next()){
 System.out.println("NAME = " + rs.getString("name"));
 }
 }
 }
}

```

Save the preceding code in a file `PgqlJdbcTest.java` and compile using:

```
javac -cp "<graph-client>/lib/*" PgqlJdbcTest.java
```

The driver is also included in a regular graph server (RPM) install. For example:

```
javac -cp "/opt/oracle/graph/lib/*" PgqlJdbcTest.java
```

### 6.9.1.4.3 Security Techniques for PGQL Queries

Programs executing dynamic queries might be subject to injection attacks that could compromise integrity and functioning of the applications.

This topic presents some techniques that can be used to prevent injection attacks when building PGQL queries using string concatenation.

- [Using Bind Variables in PGQL Queries](#)
- [Verifying PGQL Identifiers](#)

#### 6.9.1.4.3.1 Using Bind Variables in PGQL Queries

Bind variables can be used in PGQL queries for better performance and increased security. Constant scalar values in PGQL queries can be replaced with bind variables. Bind variables

are denoted by a '?' (question mark). Consider the following two queries that select people who are older than a constant age value.

```
// people older than 30
SELECT v.fname AS fname, v.lname AS lname, v.age AS age
FROM MATCH (v)
WHERE v.age > 30

// people older than 40
SELECT v.fname AS fname, v.lname AS lname, v.age AS age
FROM MATCH (v)
WHERE v.age > 40
```

The SQL translations for these queries would use the constants 30 and 40 in a similar way for the age filter. The database would perform a hard parse for each of these queries. This hard parse time can often exceed the execution time for simple queries.

You could replace the constant in each query with a bind variable as follows.

```
SELECT v.fname AS fname, v.lname AS lname, v.age AS age
FROM MATCH (v)
WHERE v.age > ?
```

This will allow the SQL engine to create a generic cursor for this query, which can be reused for different age values. As a result, a hard parse is no longer required to execute this query for different age values, and the parse time for each query will be drastically reduced.

In addition, applications that use bind variables in PGQL queries are less vulnerable to injection attacks than those that use string concatenation to embed constant values in PGQL queries.

See also *Oracle Database SQL Tuning Guide* for more information on cursor sharing and bind variables.

The `PgqlPreparedStatement` interface can be used to execute queries with bind variables as shown in `PgqlExample5.java`. `PgqlPreparedStatement` provides several set methods for different value types that can be used to set values for query execution.

There are a few limitations with bind variables in PGQL. Bind variables can only be used for constant property values. That is, vertices and edges cannot be replaced with bind variables. Also, once a particular bind variable has been set to a type, it cannot be set to a different type. For example, if `setInt(1, 30)` is executed for an `PgqlPreparedStatement`, you cannot call `setString(1, "abc")` on that same `PgqlPreparedStatement`.

### Example 6-12 `PgqlExample5.java`

`PgqlExample5.java` shows how to use bind variables with a PGQL query.

```
import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlPreparedStatement;
```



```

import oracle.pg.rdbms.pgql.PgqlResultSet;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to use bind variables with a PGQL query.
 */
public class PgqlExample5
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlPreparedStatement pps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port +" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Create a Pgql connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Query string with a bind variable (denoted by ?)
 String pgql =
 "SELECT v.\"fname\" AS fname, v.\"lname\" AS lname, v.\"age\" AS age
"+
 "FROM MATCH (v) "+
 "WHERE v.\"age\" > ?";

 // Create a PgqlPreparedStatement
 pps = pgqlConn.prepareStatement(pgql);

 // Set filter value to 30
 pps.setInt(1, 30);

 // execute query
 rs = pps.executeQuery();

```

```

 // Print query results
 System.out.println("-- Values for v.\"age\" > 30 --");
 rs.print();
 // close result set
 rs.close();

 // set filter value to 40
 pps.setInt(1, 40);

 // execute query
 rs = pps.executeQuery();

 // Print query results
 System.out.println("-- Values for v.\"age\" > 40 --");
 rs.print();
 // close result set
 rs.close();
}
finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (pps != null) {
 pps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
}
}

```

`PgqlExample5.java` has the following output for `test_graph` (which can be loaded using `GraphLoaderExample.java` code).

```

-- Values for v.age > 30 --
+-----+
| fname | lname | age |
+-----+
Susan	Blue	35
Bill	Brown	40
Ray	Green	41
+-----+		
-- Values for v.age > 40 --		
+-----+		
fname	lname	age
+-----+		
Ray	Green	41
+-----+

```

**Example 6-13 PgqlExample6.java**

PgqlExample6.java shows a query with two bind variables: one String variable and one Timestamp variable.

```
import java.sql.Connection;
import java.sql.Timestamp;

import java.time.OffsetDateTime;
import java.time.ZoneOffset;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlPreparedStatement;
import oracle.pg.rdbms.pgql.PgqlResultSet;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to use multiple bind variables with a PGQL query.
 */
public class PgqlExample6
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlPreparedStatement pps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port +" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Create a Pgql connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Query string with multiple bind variables
 String pgql =
 "SELECT v1.\"fname\" AS fname1, v2.\"fname\" AS fname2 "+
```

```

 "FROM MATCH (v1)-[e:\"knows\"]->(v2) "+
 "WHERE e.\"since\" < ? AND e.\"firstMetIn\" = ?";

// Create a PgqlPreparedStatement
pps = pgqlConn.prepareStatement(pgql);

// Set e.since < 2006-01-01T12:00:00.00Z
Timestamp t =
Timestamp.valueOf(OffsetDateTime.parse("2006-01-01T12:00:01.00Z").atZoneSameInstant(ZoneOffset.UTC).toLocalDateTime());
pps.setTimestamp(1, t);
// Set e.firstMetIn = 'CA'
pps.setString(2, "CA");

// execute query
rs = pps.executeQuery();

// Print query results
System.out.println("-- Values for e.\"since\" <
2006-01-01T12:00:01.00Z AND e.\"firstMetIn\" = 'CA' --");
rs.print();
// close result set
rs.close();

// Set e.since < 2000-01-01T12:00:00.00Z
t =
Timestamp.valueOf(OffsetDateTime.parse("2000-01-01T12:00:00.00Z").atZoneSameInstant(ZoneOffset.UTC).toLocalDateTime());
pps.setTimestamp(1, t);
// Set e.firstMetIn = 'TX'
pps.setString(2, "TX");

// execute query
rs = pps.executeQuery();

// Print query results
System.out.println("-- Values for e.\"since\" <
2000-01-01T12:00:00.00Z AND e.\"firstMetIn\" = 'TX' --");
rs.print();
// close result set
rs.close();
}
finally {
// close the result set
if (rs != null) {
 rs.close();
}
// close the statement
if (pps != null) {
 pps.close();
}
// close the connection
if (conn != null) {
 conn.close();
}
}

```

```

 }
 }
}

```

PgqlExample6.java gives the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```

-- Values for e."since" < 2006-01-01T12:00:01.00Z AND e."firstMetIn" = 'CA' --
+-----+
| FNAME1 | FNAME2 |
+-----+
| Susan | Bill |
| Susan | Ray |
+-----+
-- Values for e."since" < 2000-01-01T12:00:00.00Z AND e."firstMetIn" = 'TX' --
+-----+
| FNAME1 | FNAME2 |
+-----+
| Ray | Bill |
+-----+

```

#### 6.9.1.4.3.2 Verifying PGQL Identifiers

For some parts of a PGQL query the parser does not allow use of bind variables. In such cases, the input can be verified using the `printIdentifier` method in package `oracle.pgql.lang.ir.PgqlUtils`.

Consider the following query execution that concatenates the graph against which the graph pattern will be matched:

```
stmt.executeQuery("SELECT n.name FROM MATCH (n) ON " + graphName, "");
```

In order to avoid injection, the identifier `graphName` should be verified as follows:

```
stmt.executeQuery("SELECT n.name FROM MATCH (n) ON " +
PgqlUtils.printIdentifier(graphName), "");
```

#### 6.9.1.4.4 Using a Text Index with PGQL Queries

PGQL queries executed against Oracle Database can use Oracle Text indexes created for vertex and edge properties. After creating a text index, you can use the `CONTAINS` operator to perform a full text search. `CONTAINS` has two arguments: a vertex or edge property, and an Oracle Text search string. Any valid Oracle Text search string can be used, including advanced features such as wildcards, stemming, and soundex.

##### Example 6-14 PgqlExample7.java

PgqlExample7.java shows how to execute a `CONTAINS` query.

```
import java.sql.CallableStatement;
import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;
```

```

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to use an Oracle Text index with a PGQL
 query.
 */
public class PgqlExample7
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+":"+port+" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Create text index with SQL API
 CallableStatement cs = null;
 // text index on vertices
 cs = conn.prepareCall(
 "begin opg_apis.create_vertices_text_idx(:1,:2); end;"
);
 cs.setString(1,user);
 cs.setString(2,graph);
 cs.execute();
 cs.close();
 // text index on edges
 cs = conn.prepareCall(
 "begin opg_apis.create_edges_text_idx(:1,:2); end;"
);
 cs.setString(1,user);
 cs.setString(2,graph);
 cs.execute();
 cs.close();
 }
 }
}

```

```
// Get a PGQL connection
PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
pgqlConn.setGraph(graph);

// Create a PgqlStatement
ps = pgqlConn.createStatement();

// Query using CONTAINS text search operator on vertex property
// Find all vertices with an lname property value that starts with 'B'
String pgql =
 "SELECT v.\"fname\" AS fname, v.\"lname\" AS lname "+
 "FROM MATCH (v) "+
 "WHERE CONTAINS(v.\"lname\",'B%')";

// execute query
rs = ps.executeQuery(pgql, "");

// print results
System.out.println("-- Vertex Property Query --");
rs.print();

// close result set
rs.close();

// Query using CONTAINS text search operator on edge property
// Find all knows edges with a firstMetIn property value that ends
with 'A'
pgql =
 "SELECT v1.\"fname\" AS fname1, v2.\"fname\" AS fname2,
e.\"firstMetIn\" AS loc "+
 "FROM MATCH (v1)-[e:\"knows\"]->(v2) "+
 "WHERE CONTAINS(e.\"firstMetIn\",'%A')";

// execute query
rs = ps.executeQuery(pgql, "");

// print results
System.out.println("-- Edge Property Query --");
rs.print();

}
finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
```

```

 }
}

```

PgqlExample7.java has the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```

-- Vertex Property Query --
+-----+
| FNAME | LNAME |
+-----+
Susan	Blue
Bill	Brown
John	Black
+-----+	
-- Edge Property Query --	
+-----+	
FNAME1	FNAME1
+-----+	
Susan	Bill
John	Bill
Susan	John
Susan	Ray
+-----+

```

#### 6.9.1.4.5 Obtaining the SQL Translation for a PGQL Query

You can obtain the SQL translation for a PGQL query through methods in `PgqlStatement` and `PgqlPreparedStatement`. The raw SQL for a PGQL query can be useful for several reasons:

- You can execute the SQL directly against the database with other SQL-based tools or interfaces (for example, SQL\*Plus or SQL Developer).
- You can customize and tune the generated SQL to optimize performance or to satisfy a particular requirement of your application.
- You can build a larger SQL query that joins a PGQL subquery with other data stored in Oracle Database (such as relational tables, spatial data, and JSON data).

##### Example 6-15 PgqlExample8.java

PgqlExample8.java shows how to obtain the raw SQL translation for a PGQL query. The `translateQuery` method of `PgqlStatement` returns an `PgqlSqlQueryTrans` object that contains information about return columns from the query and the SQL translation itself.

The translated SQL returns different columns depending on the type of "logical" object or value projected from the PGQL query. A vertex or edge projected in PGQL has two corresponding columns projected in the translated SQL:

- `$IT` : id type – `NVARCHAR(1)`: 'V' for vertex or 'E' for edge
- `$ID` : vertex or edge identifier – `NUMBER`: same content as `VID` or `EID` columns in `VT$` and `GE$` tables

A property value or constant scalar value projected in PGQL has four corresponding columns projected in the translated SQL:

- `$T` : value type – `NUMBER`: same content as `T` column in `VT$` and `GE$` tables



- \$V: value – NVARCHAR2(15000): same content as V column in VT\$ and GE\$ tables
- \$VN: number value – NUMBER: same content as VN column in VT\$ and GE\$ tables
- \$VT: temporal value – TIMESTAMP WITH TIME ZONE: same content as VT column in VT\$ and GE\$ tables

```
import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlColumnDescriptor;
import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlStatement;
import oracle.pg.rdbms.pgql.PgqlSqlQueryTrans;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to obtain the SQL translation for a PGQL query.
 */
public class PgqlExample8
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port +"@"+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Create a Pgql connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // PGQL query to be translated
 String pgql =
 "SELECT v1, v1.\"fname\" AS fname1, e, e.\"since\" AS since "+
 "FROM MATCH (v1)-[e:\"knows\"]->(v2)";

 // Create a PgqlStatement
```

```
ps = pgqlConn.createStatement();

// Get the SQL translation
PgqlSqlQueryTrans sqlTrans = ps.translateQuery(pgql, "");

// Get the return column descriptions
PgqlColumnDescriptor[] cols = sqlTrans.getReturnTypes();

// Print column descriptions
System.out.println("-- Return Columns -----");
printReturnCols(cols);

// Print SQL translation
System.out.println("-- SQL Translation -----");
System.out.println(sqlTrans.getSqlTranslation());
}
finally {
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}

/**
 * Prints return columns for a SQL translation
 */
static void printReturnCols(PgqlColumnDescriptor[] cols) throws
Exception
{
 StringBuffer buff = new StringBuffer("");

 for (int i = 0; i < cols.length; i++) {

 String colName = cols[i].getColName();
 PgqlColumnDescriptor.Type colType = cols[i].getColType();
 int offset = cols[i].getSqlOffset();

 String readableType = "";
 switch(colType) {
 case VERTEX:
 readableType = "VERTEX";
 break;
 case EDGE:
 readableType = "EDGE";
 break;
 case VALUE:
 readableType = "VALUE";
 break;
 }
 }
}
```

```

 buff.append("colName="+colName+" colType="+readableType+"
offset="+offset+"\n");
 }
 System.out.println(buff.toString());
}
}
}

```

PgqlExample8.java has the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```

-- Return Columns -----
colName=[v1] colType=[VERTEX] offset=[1]
colName=[fname1] colType=[VALUE] offset=[3]
colName=[e] colType=[EDGE] offset=[7]
colName=[since] colType=[VALUE] offset=[9]
-- SQL Translation -----
SELECT n'V' AS "V1$IT",
TO$0.SVID AS "V1$ID",
TO$1.T AS "FNAME1$T",
TO$1.V AS "FNAME1$V",
TO$1.VN AS "FNAME1$VN",
TO$1.VT AS "FNAME1$VT",
n'E' AS "E$IT",
TO$0.EID AS "E$ID",
TO$0.T AS "SINCE$T",
TO$0.V AS "SINCE$V",
TO$0.VN AS "SINCE$VN",
TO$0.VT AS "SINCE$VT"
FROM (SELECT L.EID, L.SVID, L.DVID, L.EL, R.K, R.T, R.V, R.VN, R.VT
 FROM "SCOTT".TEST_GRAPHGT$ L,
 (SELECT * FROM "SCOTT".TEST_GRAPHGE$ WHERE K=n'since') R
 WHERE L.EID = R.EID(+)
) TO$0,
(SELECT L.VID, L.VL, R.K, R.T, R.V, R.VN, R.VT
 FROM "SCOTT".TEST_GRAPHVD$ L,
 (SELECT * FROM "SCOTT".TEST_GRAPHVT$ WHERE K=n'fname') R
 WHERE L.VID = R.VID(+)
) TO$1
WHERE TO$0.SVID=TO$1.VID AND
(TO$0.EL = n'knows' AND TO$0.EL IS NOT NULL)

```

### Example 6-16 PgqlExample9.java

You can also obtain the SQL translation for PGQL queries with bind variables. In this case, the corresponding SQL translation will also contain bind variables. The PgqlSqlQueryTrans interface has a getSqlBvList method that returns an ordered List of Java Objects that should be bound to the SQL query (the first Object on the list should be set at position 1, and the second should be set at position 2, and so on).

PgqlExample9.java shows how to get and execute the SQL for a PGQL query with bind variables.

```

import java.sql.Connection;
import java.sql.PreparedStatement;
import java.sql.ResultSet;
import java.sql.Timestamp;

import java.util.List;

```

```

import oracle.pg.rdbms.pgql.PgqlColumnDescriptor;
import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlPreparedStatement;
import oracle.pg.rdbms.pgql.PgqlSqlQueryTrans;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to obtain and execute the SQL translation
 for a
 * PGQL query that uses bind variables.
 */
public class PgqlExample9
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlPreparedStatement pgqlPs = null;

 PreparedStatement sqlPs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port +" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Create a Pgql connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Execute query to get a ResultSet object
 String pgql =
 "SELECT v1, v1.\"fname\" AS fname1, v1.\"age\" AS age, ? as
constVal "+
 "FROM MATCH (v1) "+
 "WHERE v1.\"fname\" = ? OR v1.\"age\" < ?";

```

```
// Create a PgqlStatement
pgqlPs = pgqlConn.prepareStatement(pgql);

// set bind values
pgqlPs.setDouble(1, 2.05d);
pgqlPs.setString(2, "Bill");
pgqlPs.setInt(3, 35);

// Get the SQL translation
PgqlSqlQueryTrans sqlTrans = pgqlPs.translateQuery("");

// Get the SQL String
String sqlStr = sqlTrans.getSqlTranslation();

// Get the return column descriptions
PgqlColumnDescriptor[] cols = sqlTrans.getReturnTypes();

// Get the bind values
List<Object> bindVals = sqlTrans.getSqlBvList();

// Print column descriptions
System.out.println("-- Return Columns -----");
printReturnCols(cols);

// Print SQL translation
System.out.println("-- SQL Translation -----");
System.out.println(sqlStr);

// Print Bind Values
System.out.println("\n-- Bind Values -----");
for (Object obj : bindVals) {
 System.out.println(obj.toString());
}

// Execute Query
// Get PreparedStatement
sqlPs = conn.prepareStatement("SELECT COUNT(*) FROM (" + sqlStr + ")");
// Set bind values and execute the PreparedStatement
executePs(sqlPs, bindVals);

// Set new bind values in the PGQL PreparedStatement
pgqlPs.setDouble(1, 3.02d);
pgqlPs.setString(2, "Ray");
pgqlPs.setInt(3, 30);

// Print Bind Values
bindVals = sqlTrans.getSqlBvList();
System.out.println("\n-- Bind Values -----");
for (Object obj : bindVals) {
 System.out.println(obj.toString());
}

// Execute the PreparedStatement with new bind values
executePs(sqlPs, bindVals);
}
```

```

finally {
 // close the SQL statement
 if (sqlPs != null) {
 sqlPs.close();
 }
 // close the statement
 if (pgqlPs != null) {
 pgqlPs.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}

/**
 * Executes a SQL PreparedStatement with the input bind values
 */
static void executePs(PreparedStatement ps, List<Object> bindVals)
throws Exception
{
 ResultSet rs = null;
 try {
 // Set bind values
 for (int idx = 0; idx < bindVals.size(); idx++) {
 Object o = bindVals.get(idx);
 // String
 if (o instanceof java.lang.String) {
 ps.setNString(idx + 1, (String)o);
 }
 // Int
 else if (o instanceof java.lang.Integer) {
 ps.setInt(idx + 1, ((Integer)o).intValue());
 }
 // Long
 else if (o instanceof java.lang.Long) {
 ps.setLong(idx + 1, ((Long)o).longValue());
 }
 // Float
 else if (o instanceof java.lang.Float) {
 ps.setFloat(idx + 1, ((Float)o).floatValue());
 }
 // Double
 else if (o instanceof java.lang.Double) {
 ps.setDouble(idx + 1, ((Double)o).doubleValue());
 }
 // Timestamp
 else if (o instanceof java.sql.Timestamp) {
 ps.setTimestamp(idx + 1, (Timestamp)o);
 }
 else {
 ps.setString(idx + 1, bindVals.get(idx).toString());
 }
 }
 }
}

```

```
 // Execute query
 rs = ps.executeQuery();
 if (rs.next()) {
 System.out.println("\n-- Execute Query: Result has "+rs.getInt(1)+"
rows --");
 }
 }
 finally {
 // close the SQL ResultSet
 if (rs != null) {
 rs.close();
 }
 }
}

/**
 * Prints return columns for a SQL translation
 */
static void printReturnCols(PgqlColumnDescriptor[] cols) throws Exception
{
 StringBuffer buff = new StringBuffer("");

 for (int i = 0; i < cols.length; i++) {

 String colName = cols[i].getColName();
 PgqlColumnDescriptor.Type colType = cols[i].getColType();
 int offset = cols[i].getSqlOffset();

 String readableType = "";
 switch(colType) {
 case VERTEX:
 readableType = "VERTEX";
 break;
 case EDGE:
 readableType = "EDGE";
 break;
 case VALUE:
 readableType = "VALUE";
 break;
 }

 buff.append("colName=["+colName+"] colType=["+readableType+"]
offset=["+offset+"]\n");
 }
 System.out.println(buff.toString());
}
}
```

PgqlExample9.java has the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```
--- Return Columns -----
colName=[v1] colType=[VERTEX] offset=[1]
colName=[fname1] colType=[VALUE] offset=[3]
```

```

colName=[age] colType=[VALUE] offset=[7]
colName=[constVal] colType=[VALUE] offset=[11]
-- SQL Translation -----
SELECT n'V' AS "V1$IT",
TO$0.VID AS "V1$ID",
TO$0.T AS "FNAME1$T",
TO$0.V AS "FNAME1$V",
TO$0.VN AS "FNAME1$VN",
TO$0.VT AS "FNAME1$VT",
TO$1.T AS "AGE$T",
TO$1.V AS "AGE$V",
TO$1.VN AS "AGE$VN",
TO$1.VT AS "AGE$VT",
4 AS "CONSTVAL$T",
to_nchar(?, 'TM9', 'NLS_Numeric_Characters='.', ''') AS "CONSTVAL$V",
? AS "CONSTVAL$VN",
to_timestamp_tz(null) AS "CONSTVAL$VT"
FROM (SELECT L.VID, L.VL, R.K, R.T, R.V, R.VN, R.VT
 FROM "SCOTT".TEST_GRAPHVD$ L,
 (SELECT * FROM "SCOTT".TEST_GRAPHVT$ WHERE K=n'fname') R
 WHERE L.VID = R.VID(+)
) TO$0,
(SELECT L.VID, L.VL, R.K, R.T, R.V, R.VN, R.VT
 FROM "SCOTT".TEST_GRAPHVD$ L,
 (SELECT * FROM "SCOTT".TEST_GRAPHVT$ WHERE K=n'age') R
 WHERE L.VID = R.VID(+)
) TO$1
WHERE TO$0.VID=TO$1.VID AND
((TO$0.T = 1 AND TO$0.V = ?) OR TO$1.VN < ?)

-- Bind Values -----
2.05
2.05
Bill
35
-- Execute Query: Result has 2 rows --

-- Bind Values -----
3.02
3.02
Ray
30
-- Execute Query: Result has 1 rows --

```

#### 6.9.1.4.6 Additional Options for PGQL Translation and Execution

Several options are available to influence PGQL query translation and execution. The following are the main ways to set query options:

- Through explicit arguments to `executeQuery` and `translateQuery`
- Through flags in the `options` string argument of `executeQuery` and `translateQuery`
- Through Java JVM arguments.

The following table summarizes the available query arguments for PGQL translation and execution.



**Table 6-4 PGQL Translation and Execution Options**

| Option                        | Default   | Explicit Argument | Options Flag       | JVM Argument                                |
|-------------------------------|-----------|-------------------|--------------------|---------------------------------------------|
| Degree of parallelism         | 0         | parallel          | none               | none                                        |
| Timeout                       | unlimited | timeout           | none               | none                                        |
| Dynamic sampling              | 2         | dynamicSampling   | none               | none                                        |
| Maximum number of results     | unlimited | maxResults        | none               | none                                        |
| GT\$ table usage              | on        | none              | USE_GT_TAB=F       | - Doracle.pg.rdbms.pgql.useGtTab=false      |
| CONNECT BY usage              | off       | none              | USE_RW=F           | -Doracle.pg.rdbms.pgql.useRW=false          |
| Distinct recursive WITH usage | off       | none              | USE_DIST_RW=T      | - Doracle.pg.rdbms.pgql.useDistRW=true      |
| Maximum path length           | unlimited | none              | MAX_PATH_LEN=n     | -Doracle.pg.rdbms.pgql.maxPathLen=n         |
| Set partial                   | false     | none              | EDGE_SET_PARTIAL=T | - Doracle.pg.rdbms.pgql.edgeSetPartial=true |
| Project null properties       | true      | none              | PROJ_NULL_PROPS=F  | - Doracle.pg.rdbms.pgql.projNullProps=false |
| VT\$ VL column usage          | on        | none              | USE_VL_COL=F       | - Doracle.pg.rdbms.pgql.useVLCol=false      |

- [Query Options Controlled by Explicit Arguments](#)
- [Using the GT\\$ Skeleton Table](#)
- [Path Query Options](#)
- [Options for Partial Object Construction](#)

#### 6.9.1.4.6.1 Query Options Controlled by Explicit Arguments

Some query options are controlled by explicit arguments to methods in the Java API.

- The `executeQuery` method of `PgqlStatement` has explicit arguments for timeout in seconds, degree of parallelism, optimizer dynamic sampling, and maximum number of results.
- The `translateQuery` method has explicit arguments for degree of parallelism, optimizer dynamic sampling, and maximum number of results. `PgqlPreparedStatement` also provides those same additional arguments for `executeQuery` and `translateQuery`.

**Example 6-17 PgqlExample10.java**

PgqlExample10.java shows PGQL query execution with additional options controlled by explicit arguments.

```
import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to execute a PGQL query with various options.
 */
public class PgqlExample10
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port+" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Execute query to get a ResultSet object
 String pgql =
 "SELECT v1.\"fname\" AS fname1, v2.\"fname\" AS fname2 "+
```

```

 "FROM MATCH (v1)-[:\"friendOf\"]->(v2)";
 rs = ps.executeQuery(pgql /* query string */,
 100 /* timeout (sec): 0 is default and implies
no timeout */,
 2 /* parallel: 1 is default */,
 6 /* dynamic sampling: 2 is default */,
 50 /* max results: -1 is default and implies no
limit */,
 "" /* options */);

 // Print query results
 rs.print();
}
finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
}
}
}

```

PgqlExample10.java gives the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```

+-----+
| FNAME1 | FNAME2 |
+-----+
Ray	Susan
John	Susan
Bill	John
Susan	John
John	Bill
+-----+

```

#### 6.9.1.4.6.2 Using the GT\$ Skeleton Table

The property graph relational schema defines a GT\$ skeleton table that stores a single row for each edge in the graph, no matter how many properties an edge has. This skeleton table is populated by default so that PGQL query execution can take advantage of the GT\$ table and avoid sorting operations on the GE\$ table in many cases, which gives a significant performance improvement.

You can add "USE\_GT\_TAB=F" to the options argument of executeQuery and translateQuery or use -Doracle.pg.rdbms.pgql.useGtTab=false in the Java command line to turn off GT\$ table usage.

**Example 6-18 PgqlExample11.java**

PgqlExample11.java shows a query that uses the GT\$ skeleton table.

```
import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlSqlQueryTrans;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to avoid using the GT$ skeleton table for
 * PGQL query execution.
 */
public class PgqlExample11
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@" + host + ":" + port + "@" + sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Execute query to get a ResultSet object
 String pgql =
 "SELECT id(v1), id(v2) "+
 "FROM MATCH (v1)-[knows]->(v2)";
 }
 }
}
```

```

// Get the SQL translation with GT table
PgqlSqlQueryTrans sqlTrans = ps.translateQuery(pgql, "");

// Print SQL translation
System.out.println("-- SQL Translation with GT Table
-----");
System.out.println(sqlTrans.getSqlTranslation());

// Get the SQL translation without GT table
sqlTrans = ps.translateQuery(pgql, "USE_GT_TAB=F");

// Print SQL translation
System.out.println("-- SQL Translation without GT Table
-----");
System.out.println(sqlTrans.getSqlTranslation());

}
finally {
// close the statement
if (ps != null) {
ps.close();
}
// close the connection
if (conn != null) {
conn.close();
}
}
}
}
}
}

```

PgqlExample11.java gives the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```

-- SQL Translation with GT Table -----
SELECT 7 AS "id(v1)$T",
to_nchar(T0$0.SVID,'TM9','NLS_Numeric_Characters='.', ''') AS "id(v1)$V",
T0$0.SVID AS "id(v1)$VN",
to_timestamp_tz(null) AS "id(v1)$VT",
7 AS "id(v2)$T",
to_nchar(T0$0.DVID,'TM9','NLS_Numeric_Characters='.', ''') AS "id(v2)$V",
T0$0.DVID AS "id(v2)$VN",
to_timestamp_tz(null) AS "id(v2)$VT"
FROM "SCOTT".TEST_GRAPHGT$ T0$0
-- SQL Translation without GT Table -----
SELECT 7 AS "id(v1)$T",
to_nchar(T0$0.SVID,'TM9','NLS_Numeric_Characters='.', ''') AS "id(v1)$V",
T0$0.SVID AS "id(v1)$VN",
to_timestamp_tz(null) AS "id(v1)$VT",
7 AS "id(v2)$T",
to_nchar(T0$0.DVID,'TM9','NLS_Numeric_Characters='.', ''') AS "id(v2)$V",
T0$0.DVID AS "id(v2)$VN",
to_timestamp_tz(null) AS "id(v2)$VT"
FROM (SELECT DISTINCT EID, SVID, DVID,EL FROM "SCOTT".TEST_GRAPHGE$) T0$0

```

### 6.9.1.4.6.3 Path Query Options

A few options are available for executing path queries in PGQL. There are two basic evaluation methods available in Oracle SQL: CONNECT BY or recursive WITH clauses. Recursive WITH is the default evaluation method. In addition, you can further modify the recursive WITH evaluation method to include a DISTINCT modifier during the recursive step of query evaluation. Computing distinct vertices at each step helps prevent a combinatorial explosion in highly connected graphs. The DISTINCT modifier is not added by default because it requires a specific parameter setting in the database ("\_recursive\_with\_control"=8).

You can also control the maximum length of paths searched. Path length in this case is defined as the number of repetitions allowed when evaluating the \* and + operators. The default maximum length is unlimited.

Path evaluation options are summarized as follows.

- **CONNECT BY:** To use CONNECT BY, specify 'USE\_RW=F' in the options argument or specify `-Doracle.pg.rdbms.pgql.useRW=false` in the Java command line.
- **Distinct Modifier in Recursive WITH:** To use the DISTINCT modifier in the recursive step, first set "\_recursive\_with\_control"=8 in your database session, then specify 'USE\_DIST\_RW=T' in the options argument or specify `-Doracle.pg.rdbms.pgql.useDistRW=true` in the Java command line. You will encounter ORA-32486: unsupported operation in recursive branch of recursive WITH clause if "\_recursive\_with\_control" has not been set to 8 in your session.
- **Path Length Restriction:** To limit maximum number of repetitions when evaluating \* and + to n, specify 'MAX\_PATH\_LEN=n' in the query options argument or specify `-Doracle.pg.rdbms.pgql.maxPathLen=n` in the Java command line.

#### Example 6-19 PgqlExample12.java

PgqlExample12.java shows path query translations under various options.

```
import java.sql.Connection;
import java.sql.Statement;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlSqlQueryTrans;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to use various options with PGQL path
 * queries.
 */
public class PgqlExample12
{
 public static void main(String[] args) throws Exception
 {
 int idx=0;
```

```

String host = args[idx++];
String port = args[idx++];
String sid = args[idx++];
String user = args[idx++];
String password = args[idx++];
String graph = args[idx++];

Connection conn = null;
PgqlStatement ps = null;

try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port +" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Set "_recursive_with_control"=8 to enable distinct optimization
 // optimization for recursive with
 Statement stmt = conn.createStatement();
 stmt.executeUpdate("alter session set \"_recursive_with_control\"=8");
 stmt.close();

 // Path Query to illustrate options
 String pgql =
 "PATH fof AS ()-[:\"friendOf\"]->() "+
 "SELECT id(v1), id(v2) "+
 "FROM MATCH (v1)-/:fof*/->(v2) "+
 "WHERE id(v1) = 2";

 // get SQL translation with defaults - Non-distinct Recursive WITH
 PgqlSqlQueryTrans sqlTrans =
 ps.translateQuery(pgql /* query string */,
 2 /* parallel: default is 1 */,
 2 /* dynamic sampling: default is 2 */,
 -1 /* max results: -1 implies no limit */,
 "" /* options */);
 System.out.println("-- Default Path Translation -----");
 System.out.println(sqlTrans.getSqlTranslation()+"\n");

 // get SQL translation with DISTINCT reachability optimization
 sqlTrans =
 ps.translateQuery(pgql /* query string */,
 2 /* parallel: default is 1 */,

```

```

 2 /* dynamic sampling: default is 2 */,
 -1 /* max results: -1 implies no limit */,
 " USE_DIST_RW=T " /* options */);
 System.out.println("-- DISTINCT RW Path Translation
-----");
 System.out.println(sqlTrans.getSqlTranslation()+"\n");

 // get SQL translation with CONNECT BY
 sqlTrans =
 ps.translateQuery(pgql /* query string */,
 2 /* parallel: default is 1 */,
 2 /* dynamic sampling: default is 2 */,
 -1 /* max results: -1 implies no limit */,
 " USE_RW=F " /* options */);
 System.out.println("-- CONNECT BY Path Translation
-----");
 System.out.println(sqlTrans.getSqlTranslation()+"\n");
 }
 finally {
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
 }
}

```

PgqlExample12.java gives the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```

-- Default Path Translation -----
SELECT /*+ parallel(2) */ * FROM(SELECT 7 AS "id(v1)$T",
to_nchar(TO$0.SVID,'TM9','NLS_Numeric_Characters='.','') AS "id(v1)$V",
TO$0.SVID AS "id(v1)$VN",
to_timestamp_tz(null) AS "id(v1)$VT",
7 AS "id(v2)$T",
to_nchar(TO$0.DVID,'TM9','NLS_Numeric_Characters='.','') AS "id(v2)$V",
TO$0.DVID AS "id(v2)$VN",
to_timestamp_tz(null) AS "id(v2)$VT"
FROM (/*Path[*/SELECT DISTINCT SVID, DVID
FROM (
SELECT 2 AS SVID, 2 AS DVID
FROM SYS.DUAL
WHERE EXISTS(
SELECT 1
FROM "SCOTT".TEST_GRAPHVTS
WHERE VID = 2)
UNION ALL
SELECT SVID,DVID FROM
(WITH RW (ROOT, DVID) AS
(SELECT ROOT, DVID FROM
(SELECT SVID ROOT, DVID
FROM (SELECT TO$0.SVID AS SVID,

```



```

T0$0.DVID AS DVID
FROM "SCOTT".TEST_GRAPHGT$ T0$0
WHERE T0$0.SVID = 2 AND
(T0$0.EL = n'friendOf' AND T0$0.EL IS NOT NULL)
) UNION ALL
SELECT RW.ROOT, R.DVID
FROM (SELECT T0$0.SVID AS SVID,
T0$0.DVID AS DVID
FROM "SCOTT".TEST_GRAPHGT$ T0$0
WHERE (T0$0.EL = n'friendOf' AND T0$0.EL IS NOT NULL)) R, RW
WHERE RW.DVID = R.SVID)
CYCLE DVID SET cycle_col TO 1 DEFAULT 0
SELECT ROOT SVID, DVID FROM RW))]Path*/) T0$0
WHERE T0$0.SVID = 2)

-- DISTINCT RW Path Translation -----
SELECT /*+ parallel(2) */ * FROM(SELECT 7 AS "id(v1)$T",
to_nchar(T0$0.SVID,'TM9','NLS_Numeric_Characters='.',') AS "id(v1)$V",
T0$0.SVID AS "id(v1)$VN",
to_timestamp_tz(null) AS "id(v1)$VT",
7 AS "id(v2)$T",
to_nchar(T0$0.DVID,'TM9','NLS_Numeric_Characters='.',') AS "id(v2)$V",
T0$0.DVID AS "id(v2)$VN",
to_timestamp_tz(null) AS "id(v2)$VT"
FROM (/*Path[*/SELECT DISTINCT SVID, DVID
FROM (
SELECT 2 AS SVID, 2 AS DVID
FROM SYS.DUAL
WHERE EXISTS(
SELECT 1
FROM "SCOTT".TEST_GRAPHVT$
WHERE VID = 2)
UNION ALL
SELECT SVID,DVID FROM
(WITH RW (ROOT, DVID) AS
(SELECT ROOT, DVID FROM
(SELECT SVID ROOT, DVID
FROM (SELECT T0$0.SVID AS SVID,
T0$0.DVID AS DVID
FROM "SCOTT".TEST_GRAPHGT$ T0$0
WHERE T0$0.SVID = 2 AND
(T0$0.EL = n'friendOf' AND T0$0.EL IS NOT NULL))
) UNION ALL
SELECT DISTINCT RW.ROOT, R.DVID
FROM (SELECT T0$0.SVID AS SVID,
T0$0.DVID AS DVID
FROM "SCOTT".TEST_GRAPHGT$ T0$0
WHERE (T0$0.EL = n'friendOf' AND T0$0.EL IS NOT NULL)) R, RW
WHERE RW.DVID = R.SVID)
CYCLE DVID SET cycle_col TO 1 DEFAULT 0
SELECT ROOT SVID, DVID FROM RW))]Path*/) T0$0
WHERE T0$0.SVID = 2)

-- CONNECT BY Path Translation -----
SELECT /*+ parallel(2) */ * FROM(SELECT 7 AS "id(v1)$T",
to_nchar(T0$0.SVID,'TM9','NLS_Numeric_Characters='.',') AS "id(v1)$V",
T0$0.SVID AS "id(v1)$VN",
to_timestamp_tz(null) AS "id(v1)$VT",
7 AS "id(v2)$T",
to_nchar(T0$0.DVID,'TM9','NLS_Numeric_Characters='.',') AS "id(v2)$V",
T0$0.DVID AS "id(v2)$VN",

```

```

to_timestamp_tz(null) AS "id(v2)$VT"
FROM (/*Path[*/SELECT DISTINCT SVID, DVID
FROM (
SELECT 2 AS SVID, 2 AS DVID
FROM SYS.DUAL
WHERE EXISTS(
SELECT 1
FROM "SCOTT".TEST_GRAPHVT$
WHERE VID = 2)
UNION ALL
SELECT SVID, DVID
FROM
(SELECT CONNECT_BY_ROOT T0$0.SVID AS SVID, T0$0.DVID AS DVID
FROM(
SELECT T0$0.SVID AS SVID,
T0$0.DVID AS DVID
FROM "SCOTT".TEST_GRAPHGT$ T0$0
WHERE (T0$0.EL = n'friendOf' AND T0$0.EL IS NOT NULL)) T0$0
START WITH T0$0.SVID = 2
CONNECT BY NOCYCLE PRIOR DVID = SVID))/*]Path*/) T0$0
WHERE T0$0.SVID = 2)

```

The query plan for the first query with the default recursive WITH strategy should look similar to the following.

```

-- default RW

| Id | Operation |
Name
0
1
2
SYS_TEMP_0FD9D6662_37AA44
3
4
5
:TQ20000
6
SYS_TEMP_0FD9D6662_37AA44
7
* 8
TEST_GRAPHGT$
* 9
TEST_GRAPHXSG$
10
11
:TQ10000
12
SYS_TEMP_0FD9D6662_37AA44
13
14

```

```

| |
* 15	TABLE ACCESS FULL	SYS_TEMP_0FD9D6662_37AA44
16	PARTITION HASH ALL	
* 17	TABLE ACCESS BY LOCAL INDEX ROWID BATCHED	TEST_GRAPHGT$
* 18	INDEX RANGE SCAN	TEST_GRAPHXSG$
19	PX COORDINATOR	
20	PX SEND QC (RANDOM)	:TQ30001
21	VIEW	
22	HASH UNIQUE	
23	PX RECEIVE	
24	PX SEND HASH	:TQ30000
25	HASH UNIQUE	
26	VIEW	
27	UNION-ALL	
28	PX SELECTOR	
* 29	FILTER	
30	FAST DUAL	
31	PARTITION HASH SINGLE	
* 32	INDEX SKIP SCAN	TEST_GRAPHXQV$
33	VIEW	
* 34	VIEW	
35	PX BLOCK ITERATOR	
36	TABLE ACCESS FULL	SYS_TEMP_0FD9D6662_37AA44

```

The query plan for the second query that adds a DISTINCT modifier in the recursive step should look similar to the following.

```

| Id | Operation |
| Name | |

0	SELECT STATEMENT
1	TEMP TABLE TRANSFORMATION
2	LOAD AS SELECT (CURSOR DURATION MEMORY)
SYS_TEMP_0FD9D6669_37AA44	
3	UNION ALL (RECURSIVE WITH) BREADTH FIRST
4	PX COORDINATOR
5	PX SEND QC (RANDOM)
:TQ20000	
6	LOAD AS SELECT (CURSOR DURATION MEMORY)
SYS_TEMP_0FD9D6669_37AA44	
7	PX PARTITION HASH ALL
* 8	TABLE ACCESS BY LOCAL INDEX ROWID BATCHED
TEST_GRAPHGT$	
* 9	INDEX RANGE SCAN
TEST_GRAPHXSG$	
10	PX COORDINATOR
11	PX SEND QC (RANDOM)
:TQ10001	
12	LOAD AS SELECT (CURSOR DURATION MEMORY)
SYS_TEMP_0FD9D6669_37AA44	
13	SORT GROUP BY

```

```

14		PX RECEIVE
15		PX SEND HASH
:TQ10000		
16		SORT GROUP BY
17		NESTED LOOPS
18		PX BLOCK ITERATOR
* 19		TABLE ACCESS FULL
SYS_TEMP_0FD9D6669_37AA44		
20		PARTITION HASH ALL
* 21		TABLE ACCESS BY LOCAL INDEX ROWID BATCHED
TEST_GRAPHGT$		
* 22		INDEX RANGE SCAN
TEST_GRAPHXSG$		
23		PX COORDINATOR
24		PX SEND QC (RANDOM)
:TQ30001		
25		VIEW
26		HASH UNIQUE
27		PX RECEIVE
28		PX SEND HASH
:TQ30000		
29		HASH UNIQUE
30		VIEW
31		UNION-ALL
32		PX SELECTOR
* 33		FILTER
34		FAST DUAL
35		PARTITION HASH SINGLE
* 36		INDEX SKIP SCAN
TEST_GRAPHXQV$		
37		VIEW
* 38		VIEW
39		PX BLOCK ITERATOR
40		TABLE ACCESS FULL
SYS_TEMP_0FD9D6669_37AA44 |


```

The query plan for the third query that uses CONNECTY BY should look similar to the following.

| Id   | Operation                                 | Name            |
|------|-------------------------------------------|-----------------|
| 0    | SELECT STATEMENT                          |                 |
| 1    | VIEW                                      |                 |
| 2    | HASH UNIQUE                               |                 |
| 3    | VIEW                                      |                 |
| 4    | UNION-ALL                                 |                 |
| * 5  | FILTER                                    |                 |
| 6    | FAST DUAL                                 |                 |
| 7    | PARTITION HASH SINGLE                     |                 |
| * 8  | INDEX SKIP SCAN                           | TEST_GRAPHXQV\$ |
| * 9  | VIEW                                      |                 |
| * 10 | CONNECT BY WITH FILTERING                 |                 |
| 11   | PX COORDINATOR                            |                 |
| 12   | PX SEND QC (RANDOM)                       | :TQ10000        |
| 13   | PX PARTITION HASH ALL                     |                 |
| * 14 | TABLE ACCESS BY LOCAL INDEX ROWID BATCHED | TEST_GRAPHGT\$  |
| * 15 | INDEX RANGE SCAN                          | TEST_GRAPHXSG\$ |
| 16   | NESTED LOOPS                              |                 |
| 17   | CONNECT BY PUMP                           |                 |
| 18   | PARTITION HASH ALL                        |                 |
| * 19 | TABLE ACCESS BY LOCAL INDEX ROWID BATCHED | TEST_GRAPHGT\$  |
| * 20 | INDEX RANGE SCAN                          | TEST_GRAPHXSG\$ |

**Example 6-20 PgqlExample13.java**

PgqlExample13.java shows how to set length restrictions during path query evaluation.

```
import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to use the maximum path length option for
 * PGQL path queries.
 */
public class PgqlExample13
{
 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
```

```

PgqlResultSet rs = null;

try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
pds.setURL("jdbc:oracle:thin:@"+host+": "+port +":"+sid);
pds.setUser(user);
pds.setPassword(password);
conn = pds.getConnection();

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Path Query to illustrate options
 String pgql =
 "PATH fof AS ()-[:\"friendOf\"]->() "+
 "SELECT v1.\"fname\" AS fname1, v2.\"fname\" AS fname2 "+
 "FROM MATCH (v1)-/:fof*/->(v2) "+
 "WHERE v1.\"fname\" = 'Ray'";

 // execute query for 1-hop
 rs = ps.executeQuery(pgql, " MAX_PATH_LEN=1 ");

 // print results
 System.out.println("-- Results for 1-hop -----");
 rs.print();

 // close result set
 rs.close();

 // execute query for 2-hop
 rs = ps.executeQuery(pgql, " MAX_PATH_LEN=2 ");

 // print results
 System.out.println("-- Results for 2-hop -----");
 rs.print();

 // close result set
 rs.close();

 // execute query for 3-hop
 rs = ps.executeQuery(pgql, " MAX_PATH_LEN=3 ");

 // print results
 System.out.println("-- Results for 3-hop -----");
 rs.print();
}

```

```

 // close result set
 rs.close();

 }
 finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
 }
}
}
}

```

PgqlExample13.java has the following output for test\_graph (which can be loaded using GraphLoaderExample.java code).

```

-- Results for 1-hop -----
+-----+
| FNAME1 | FNAME2 |
+-----+
| Ray | Ray |
| Ray | Susan |
+-----+
-- Results for 2-hop -----
+-----+
| FNAME1 | FNAME2 |
+-----+
Ray	Susan
Ray	Ray
Ray	John
+-----+	
-- Results for 3-hop -----	
+-----+	
FNAME1	FNAME2
+-----+	
Ray	Susan
Ray	Bill
Ray	Ray
Ray	John
+-----+

```

#### 6.9.1.4.6.4 Options for Partial Object Construction

When reading edges from a query result, there are two possible behaviors when adding the start and end vertex to any local caches:

- Add only the vertex ID, which is available from the edge itself. This option is the default, for efficiency.

- Add the vertex ID, and retrieve all properties for the start and end vertex. For this behavior, you can call `setPartial(true)` on each `OracleVertex` object constructed from your PGQL query result set.

#### Example 6-21 `PgqlExample14.java`

`PgqlExample14.java` illustrates this difference in behavior. This program first executes a query to retrieve all edges, which causes the incident vertices to be added to a local cache. The second query retrieves all vertices. The program then prints each `OracleVertex` object to show which properties have been loaded.

```
import java.sql.Connection;

import oracle.pg.rdbms.Oracle;
import oracle.pg.rdbms.OraclePropertyGraph;
import oracle.pg.rdbms.OracleVertex;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows the behavior of setPartial(true) for
 * OracleVertex objects
 * created from PGQL query results.
 */
public class PgqlExample14
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 Oracle oracle = null;
 OraclePropertyGraph opg = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@" + host + ":" + port + ":" + sid);
```



```
pds.setUser(user);
pds.setPassword(password);
conn = pds.getConnection();

// Get a PGQL connection
PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
pgqlConn.setGraph(graph);

// Create a PgqlStatement
ps = pgqlConn.createStatement();

// Query to illustrate set partial
String pgql =
 "SELECT id(e), label(e) "+
 "FROM MATCH (v1)-[e:\"knows\"]->(v2)";

// execute query
rs = ps.executeQuery(pgql, " ");

// print results
System.out.println("-- Results for edge query -----");
rs.print();

// close result set
rs.close();

// Create an Oracle Property Graph instance
oracle = new Oracle(conn);
opg = OraclePropertyGraph.getInstance(oracle, graph);

// Query to retrieve vertices
pgql =
 "SELECT id(v) "+
 "FROM MATCH (v)";

// Get each vertex object in result and print with toString()
rs = ps.executeQuery(pgql, " ");

// iterate through result
System.out.println("-- Vertex objects retrieved from vertex query --");
while (rs.next()) {
 Long vid = rs.getLong(1);
 OracleVertex v = OracleVertex.getInstance(opg, vid);
 System.out.println(v.toString());
}
// close result set
rs.close();

// Execute the same query but call setPartial(true) for each vertex
rs = ps.executeQuery(pgql, " ");
System.out.println("-- Vertex objects retrieved from vertex query with
setPartial(true) --");
while (rs.next()) {
 Long vid = rs.getLong(1);
 OracleVertex v = OracleVertex.getInstance(opg, vid);
```



```

Vertex ID 3 [NULL] {bval:bol:false, fname:str:Susan, lname:str:Blue, mval:bol:false,
age:int:35}
Vertex ID 0 [NULL] {bval:bol:true, fname:str:Bill, lname:str:Brown, mval:str:y,
age:int:40}
Vertex ID 2 [NULL] {fname:str:Ray, lname:str:Green, mval:dat:1985-01-01 04:00:00.0,
age:int:41}
Vertex ID 1 [NULL] {bval:bol:true, fname:str:John, lname:str:Black, mval:int:27,
age:int:30}

```

#### 6.9.1.4.7 Querying Another User's Property Graph

You can query another user's property graph data if you have been granted the appropriate privileges in the database. For example, to query GRAPH1 in SCOTT's schema, you must have READ privilege on SCOTT.GRAPH1GE\$, SCOTT.GRAPH1VT\$, SCOTT.GRAPH1GT\$, and SCOTT.GRAPH1VD\$.

##### Example 6-22 PgqlExample15.java

PgqlExample15.java shows how another user can query a graph in SCOTT's schema.

```

import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to query a property graph located in another user's
 * schema. READ privilege on GE$, VT$, GT$ and VD$ tables for the other
 * user's
 * property graph are required to avoid ORA-00942: table or view does not
 * exist.
 */
public class PgqlExample15
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection

```

```
PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
pds.setURL("jdbc:oracle:thin:@'+host+':" +port +":"+sid);
pds.setUser(user);
pds.setPassword(password);
conn = pds.getConnection();

// Get a PGQL connection
PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
pgqlConn.setGraph(graph);

// Set schema so that we can query Scott's graph
pgqlConn.setSchema("SCOTT");

// Create a PgqlStatement
ps = pgqlConn.createStatement();

// Execute query to get a ResultSet object
String pgql =
 "SELECT v.\"fname\" AS fname, v.\"lname\" AS lname "+
 "FROM MATCH (v)";
rs = ps.executeQuery(pgql, "");

// Print query results
rs.print();
}
finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
```

The following SQL statements create database user USER2 and grant the necessary privileges. You can also use the `OraclePropertyGraph.grantAccess` Java API to achieve the same effect.

```
SQL> grant connect, resource, unlimited tablespace to user2 identified by user2;
```

```
Grant succeeded.
```

```
SQL> grant read on scott.test_graphvt$ to user2;
```

```

Grant succeeded.

SQL> grant read on scott.test_graphge$ to user2;

Grant succeeded.

SQL> grant read on scott.test_graphgt$ to user2;

Grant succeeded.

SQL> grant read on scott.test_graphvd$ to user2;

Grant succeeded.

```

The output for `PgqlExample15.java` for the `test_graph` data set when connected to the database as `USER2` is as follows. Note that `test_graph` should have already been loaded (using `GraphLoaderExample.java` code) as `GRAPH1` by user `SCOTT` before running `PgqlExample15`.

```

+-----+
| FNAME | LNAME |
+-----+
Susan	Blue
Bill	Brown
Ray	Green
John	Black
+-----+

```

#### 6.9.1.4.8 Using Query Optimizer Hints with PGQL

The Java API allows query optimizer hints that influence the join type when executing PGQL queries. The `executeQuery` and `translateQuery` methods in `PgqlStatement` and `PgqlPreparedStatement` accept the following strings in the `options` argument to influence the query plan for the corresponding SQL query.

- `ALL_EDGE_NL` – Use Nested Loop join for all joins that involve the `$GE` and `$GT` tables.
- `ALL_EDGE_HASH` – Use HASH join for all joins that involve the `$GE` and `$GT` tables.
- `ALL_VERTEX_NL` – Use Nested Loop join for all joins that involve the `$VT` table.
- `ALL_VERTEX_HASH` – Use HASH join for all joins that involve the `$VT` table.

#### Example 6-23 `PgqlExample16.java`

`PgqlExample16.java` shows how to use optimizer hints to influence the joins used for a graph traversal.

```

import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlSqlQueryTrans;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to use query optimizer hints with PGQL queries.
 */

```

```

public class PgqlExample16
{
 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+":"+port +" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();
 // Query to illustrate join hints
 String pgql =
 "SELECT id(v1), id(v4) "+
 "FROM MATCH (v1)-[:\"friendOf\"]->(v2)-[:\"friendOf\"]-
 >(v3)-[:\"friendOf\"]->(v4)";

 // get SQL translation with hash join hint
 PgqlSqlQueryTrans sqlTrans =
 ps.translateQuery(pgql /* query string */,
 " ALL_EDGE_HASH " /* options */);
 // print SQL translation
 System.out.println("-- Query with ALL_EDGE_HASH
 -----");
 System.out.println(sqlTrans.getSqlTranslation()+"\n");

 // get SQL translation with nested loop join hint
 sqlTrans =
 ps.translateQuery(pgql /* query string */,
 " ALL_EDGE_NL " /* options */);
 // print SQL translation
 System.out.println("-- Query with ALL_EDGE_NL
 -----");
 }
 }
}

```



```

| Id | Operation | Name |

0	SELECT STATEMENT	
* 1	HASH JOIN	
* 2	HASH JOIN	
3	PARTITION HASH ALL	
* 4	TABLE ACCESS FULL	TEST_GRAPHGT$
5	PARTITION HASH ALL	
* 6	TABLE ACCESS FULL	TEST_GRAPHGT$
7	PARTITION HASH ALL	
* 8	TABLE ACCESS FULL	TEST_GRAPHGT$

```

The query plan for the second query that uses ALL\_EDGE\_NL should look similar to the following.

```

| Id | Operation | Name |

0	SELECT STATEMENT	
1	NESTED LOOPS	
2	NESTED LOOPS	
3	PARTITION HASH ALL	
* 4	TABLE ACCESS FULL	TEST_GRAPHGT$
5	PARTITION HASH ALL	
* 6	TABLE ACCESS BY LOCAL INDEX ROWID BATCHED	TEST_GRAPHGT$
* 7	INDEX RANGE SCAN	TEST_GRAPHXSG$
8	PARTITION HASH ALL	
* 9	TABLE ACCESS BY LOCAL INDEX ROWID BATCHED	TEST_GRAPHGT$
* 10	INDEX RANGE SCAN	TEST_GRAPHXSG$

```

#### 6.9.1.4.9 Modifying Property Graphs through INSERT, UPDATE, and DELETE Statements

PGQL supports INSERT, UPDATE, and DELETE operations on Property Graphs. The method `execute` in `PgqlStatement` lets you execute such DML operations. This topic provides several examples of such operations.

#### Note:

JDBC connection `autocommit` must be off in order to be able to execute INSERT, UPDATE, and DELETE statements.

#### Example 6-24 `PgqlExample17.java` (Insert)

`PgqlExample17.java` inserts several vertices and edges into a graph. Notice that the special property `_ora_id` is used to define ID values of vertices and edges. If the property `_ora_id` is omitted, a unique ID is generated for each new vertex or edge that is inserted into the graph.

```

import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;

```



```

import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to execute a PGQL INSERT operation.
 */
public class PgqlExample17
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+":"+port+":"+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Execute insert statement
 String pgql =
 "INSERT VERTEX p1 LABELS (person) PROPERTIES (p1.\"_ora_id\" = 1,
 p1.fname = 'Jake') "+
 " , VERTEX p2 LABELS (person) PROPERTIES (p2.\"_ora_id\" = 2,
 p2.fname = 'Amy') "+
 " , VERTEX p3 LABELS (person) PROPERTIES (p3.\"_ora_id\" = 3,
 p3.fname = 'Erik') "+
 " , VERTEX p4 LABELS (person) PROPERTIES (p4.\"_ora_id\" = 4,
 p4.fname = 'Jane') "+
 " , EDGE e1 BETWEEN p1 AND p2 LABELS (knows) PROPERTIES
 (e1.\"_ora_id\" = 1, e1.since = DATE '2003-04-21') "+

```



```
import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to execute a PGQL UPDATE operation.
 */
public class PgqlExample18
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@" + host + ":" + port + ":" + sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Execute update statement
 String pgql =
 "UPDATE p1 SET (p1.age = 47, p1.lname = 'Red'), "+
 " p2 SET (p2.age = 29, p2.lname = 'White'), "+
 " e SET (e.strength = 100) "+
 "FROM MATCH (p1) -[e:knows]-> (p2) "+
 "WHERE p1.fname = 'Jake' AND p2.fname = 'Amy'";
 ps.execute(pgql, /* query string */
 "", /* query options */
 "" /* modify options */);
```



```

import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to execute a PGQL DELETE operation.
 */
public class PgqlExample19
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+":"+port+" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Execute delete statement
 String pgql =
 "DELETE e "+
 " FROM MATCH (p1) -[e:knows]-> (p2) "+
 " WHERE p1.fname = 'Jake'";
 ps.execute(pgql, /* query string */
 "", /* query options */
 "" /* modify options */);

 // Execute a query to verify delete
 pgql =
 "SELECT p1.fname AS fname1, p2.fname AS fname2 "+
 " FROM MATCH (p1) -[e:knows]-> (p2)";
 rs = ps.executeQuery(pgql, "");
 }
 }
}

```

```

 // Print the results
 rs.print();
 }
 finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
 }
}
}
}

```

The output for `PgqlExample19.java` applied on a graph where `PgqlExample18.java` has been previously executed is:

```

+-----+
| FNAME1 | FNAME2 |
+-----+
| Erik | Jane |
+-----+

```

For more examples of DELETE statement, see the relevant section of the PGQL specification [here](#).

### Example 6-27 PgqlExample20.java (Multiple Modifications)

`PgqlExample20.java` executes multiple modifications in the same statement: an edge is inserted, vertex properties are updated, and another edge is deleted.

```

import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to execute a PGQL
 * INSERT/UPDATE/DELETE operation.
 */
public class PgqlExample20
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;

```

```

String host = args[idx++];
String port = args[idx++];
String sid = args[idx++];
String user = args[idx++];
String password = args[idx++];
String graph = args[idx++];

Connection conn = null;
PgqlStatement ps = null;
PgqlResultSet rs = null;

try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port+" "+sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
 ps = pgqlConn.createStatement();

 // Execute INSERT/UPDATE/DELETE statement
 String pgql =
 "INSERT EDGE f BETWEEN p2 AND p1 LABELS (knows) PROPERTIES (f.since
= e.since) "+
 "UPDATE p1 SET (p1.age = 30) "+
 " , p2 SET (p2.age = 25) "+
 "DELETE e "+
 " FROM MATCH (p1) -[e:knows]-> (p2) "+
 " WHERE p1.fname = 'Erik'";
 ps.execute(pgql, /* query string */
 "", /* query options */
 "" /* modify options */);

 // Execute a query to verify INSERT/UPDATE/DELETE
 pgql =
 "SELECT p1.fname AS fname1, p1.age AS age1, "+
 " p2.fname AS fname2, p2.age AS age2, e.since "+
 " FROM MATCH (p1) -[e:knows]-> (p2)";
 rs = ps.executeQuery(pgql, "");

 // Print the results
 rs.print();
}
finally {
 // close the result set
 if (rs != null) {

```

```

 rs.close();
 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}

```

The output for `PgqlExample20.java` applied on a graph where `PgqlExample19.java` has been previously executed is:

```

+-----+
| FNAME1 | AGE1 | FNAME2 | AGE2 | SINCE |
+-----+
| Jane | 25 | Erik | 30 | 1999-01-02 16:00:00.0 |
+-----+

```

For more examples of INSERT/UPDATE/DELETE statements, see the relevant section of the PGQL specification [here](#).

- [Additional Options for PGQL Statement Execution](#)

#### 6.9.1.4.9.1 Additional Options for PGQL Statement Execution

Several options are available to influence PGQL statement execution. The following are the main ways to set query options:

- Through flags in the `modify options` string argument of `execute`
- Through Java JVM arguments.

The following table summarizes the main options for modifying PGQL statement execution.

**Table 6-5 PGQL Statement Modification Options**

| Option         | Default                                                          | Options Flag     | JVM Argument                                   |
|----------------|------------------------------------------------------------------|------------------|------------------------------------------------|
| Auto commit    | true if JDBC auto commit is off, false if JDBC auto commit is on | AUTO_COMMIT=F    | -<br>Doracle.pg.rdbms.pgql.autoCommit=false    |
| Delete cascade | true                                                             | DELETE_CASCADE=F | -<br>Doracle.pg.rdbms.pgql.deleteCascade=false |

- [Turning Off PGQL Auto Commit](#)
- [Turning Off Cascading Deletion](#)

##### 6.9.1.4.9.1.1 Turning Off PGQL Auto Commit

When an INSERT, UPDATE, or DELETE operation is executed, a commit is performed automatically at the end of the PGQL execution so that changes are persisted on the RDBMS side.



The flag `AUTO_COMMIT=F` can be added to the `options` argument of `execute` or the flag `Doracle.pg.rdbms.pgql.autoCommit=false` can be set in the Java command line to turn off auto commit. Notice that when auto commit is off, you must perform any necessary commits or rollbacks on the JDBC connection in order to persist or cancel graph modifications.

### Example 6-28 Turn Off Auto Commit and Roll Back Changes

`PgqlExample21.java` turns off auto commit and performs a rollback of the changes.

```
import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to modify a PGQL graph
 * with auto commit off.
 */
public class PgqlExample21
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;
 PgqlResultSet rs = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();
 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@" + host + ":" + port + ":" + sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
 pgqlConn.setGraph(graph);

 // Create a PgqlStatement
```

```

ps = pgqlConn.createStatement();

// Delete all the edges in the graph
String pgql =
 "DELETE e "+
 " FROM MATCH () -[e]-> ()";
ps.execute(pgql, /* query string */
 "", /* query options */
 "AUTO_COMMIT=F" /* modify options */);

// Execute a query to verify deletion
pgql =
 "SELECT COUNT(e) "+
 " FROM MATCH () -[e]-> ()";
rs = ps.executeQuery(pgql, "");

// Print the results
System.out.println("Number of edges after deletion:");
rs.print();
rs.close();

// Rollback the changes. This is possible because
// AUTO_COMMIT=F flag was used in execute
conn.rollback();

// Execute a query to verify rollback
pgql =
 "SELECT COUNT(e) "+
 " FROM MATCH () -[e]-> ()";
rs = ps.executeQuery(pgql, "");

// Print the results
System.out.println("Number of edges after rollback:");
rs.print();
}
finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
}

```

PgqlExample21.java gives the following output for a graph with one edge:

```

Number of edges after deletion:
+-----+
| COUNT(e) |
+-----+
| 0 |
+-----+
Number of edges after rollback:
+-----+
| COUNT(e) |
+-----+
| 1 |
+-----+

```

#### 6.9.1.4.9.1.2 Turning Off Cascading Deletion

When a vertex is deleted from a graph, all its input and output edges are also deleted automatically.

Using the flag `DELETE_CASCADE=F` in the `options` argument of `execute` or setting the flag or setting the flag `Oracle.pg.rdbms.pgql.autoCommit=false` in the Java command line lets you turn off cascading deletion. When a vertex with input or output edges is deleted and cascading deletion is off, an error is thrown to warn about the unsafe operation that you are trying to perform.

#### Example 6-29 Turn Off Cascading Deletion

`PgqlExample22.java` attempts to delete a vertex with an output edge when cascading deletion is off.

```

import java.sql.Connection;

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlStatement;
import oracle.pg.rdbms.pgql.PgqlToSQLException;

import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows the use of DELETE_CASCADE flag.
 */
public class PgqlExample22
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String graph = args[idx++];

 Connection conn = null;
 PgqlStatement ps = null;

 try {

```

```

//Get a jdbc connection
PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
pds.setURL("jdbc:oracle:thin:@"+host+": "+port +":"+sid);
pds.setUser(user);
pds.setPassword(password);
conn = pds.getConnection();
conn.setAutoCommit(false);

// Get a PGQL connection
PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);
pgqlConn.setGraph(graph);

// Create a PgqlStatement
ps = pgqlConn.createStatement();

// Delete all the vertices with output edges
// This will throw an error
String pgql =
 "DELETE v "+
 " FROM MATCH (v) -[e]-> ()";
ps.execute(pgql, /* query string */
 "", /* query options */
 "DELETE_CASCADE=F" /* modify options */);
}
catch (PgqlToSQLException ex){
 System.out.println("Error in execution: " + ex.getMessage());
}
finally {
 // close the statement
 if (ps != null) {
 ps.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
}
}

```

PgqlExample22.java gives the following output for a graph with at least one edge:

```

Error in execution: Attempting to delete vertices with incoming/outgoing edges.
Drop edges first or turn on DELETE_CASCADE option

```

### 6.9.1.5 Using the Python Client to Execute PGQL Queries

You can use the new Python package `OPG4Py` for executing PGQL queries against Oracle Database. This new package contains a sub-package `Pgql` with one or more modules that wraps around the Java API in the `oracle.pg.rdbms.pgql` package.

See [Python API Reference](#) for more information.

- [Creating a Property Graph Using the Python Client](#)
- [Dropping a Property Graph Using the Python Client](#)
- [Basic Query Execution](#)
- [Iterating a Query Result Set](#)

### 6.9.1.5.1 Creating a Property Graph Using the Python Client

You can create a property graph using the `CREATE PROPERTY GRAPH` statement in Python.

#### Creating a Property Graph Using the Python Client

- Launch the Python client as shown:

```
./bin/opg4py --no_connect
```

- Create a PGQL connection to connect to the database as shown:

```
>>> pgql_conn = opg4py.pgql.get_connection(<user>, <password>, <jdbc_url>)
PgqlConnection(schema: GRAPHUSER, graph: None)
```

- Create a PGQL statement as shown:

```
>>> pgql_statement = pgql_conn.create_statement()
PgqlStatement(java_pgql_statement: oracle.pg.rdbms.pgql.PgqlStatement)
```

- Define and execute the `CREATE PROPERTY GRAPH` statement as shown:

```
pgql = """
CREATE PROPERTY GRAPH <graph_name>
VERTEX TABLES (
 bank_accounts
 LABEL accounts
 PROPERTIES ALL COLUMNS
)
EDGE TABLES (
 bank_txns
 SOURCE KEY (from_acct_id) REFERENCES bank_accounts
 DESTINATION KEY (to_acct_id) REFERENCES bank_accounts
 LABEL transfers PROPERTIES ALL COLUMNS
)
"""
```

where `<graph_name>` is the name of the graph.

```
pgql_statement.execute(pgql)
```

The graph gets created.

### 6.9.1.5.2 Dropping a Property Graph Using the Python Client

You can drop a property graph using the `DROP PROPERTY GRAPH` statement in Python.

### Dropping a Property Graph Using the Python Client

- Define and execute the `DROP PROPERTY GRAPH` statement as shown:

```
>>> pgql = "DROP PROPERTY GRAPH <graph_name>"
```

where `<graph_name>` is the name of the graph.

```
>>> pgql_statement.execute(pgql)
```

The graph gets dropped.

### 6.9.1.5.3 Basic Query Execution

You can execute PGQL queries using the `opg4py.pgql` Python wrapper.

#### Executing PGQL Queries Using the Python Client

- Set the graph for querying as shown:

```
>>> pgql_conn.set_graph("<graph_name>")
```

where `<graph_name>` is the name of the graph.

- Define and execute the PGQL `SELECT` query. For example,

```
>>> pgql = "SELECT e.from_acct_id, e.to_acct_id, e.amount FROM
MATCH (n:accounts) -[e:transfers]-> (m:accounts) on bank_graph
limit 10"
```

- Execute and print the result set as shown:

```
>>> pgql_result_set = pgql_statement.execute_query(pgql)
>>> pgql_result_set.print()
```

```

+-----+
| FROM_ACCT_ID | TO_ACCT_ID | AMOUNT |
+-----+
781.0	712.0	1000.0
190.0	555.0	1000.0
191.0	329.0	1000.0
198.0	57.0	1000.0
220.0	441.0	1000.0
251.0	387.0	1000.0
254.0	188.0	1000.0
259.0	305.0	1000.0
261.0	145.0	1000.0
263.0	40.0	1000.0
+-----+
```

```
PgqlResultSet(java_pgql_result_set:
oracle.pg.rdbms.pgql.PgqlResultSet, # of results: 0)
```

Also, you can convert the PGQL result set obtained in the preceding code to a Pandas dataframe using the `to_pandas()` method.

 **Note:**

The `pandas` package must be installed in your system to successfully execute the call to `to_pandas()`. This package is automatically installed at the time of the Python client installation. However, if your call to `to_pandas()` fails, verify if the `pandas` module is installed in your system. In case the module is found missing, then install the `pandas` package manually.

#### 6.9.1.5.4 Iterating a Query Result Set

You can iterate your query result set using the methods in `PgqlResultSet`.

You can position the cursor for iterating your query result set using the following methods:

- `first()` : boolean
- `next()` : boolean
- `previous()` : boolean
- `last()` : boolean
- `before_first()`
- `after_last()`
- `absolute(target_row_value)` : boolean
- `relative(offset_value)` : boolean

Once the cursor is positioned at the desired row, you can use the following getters to obtain values:

- `get(column_idx)` : Object
- `get(column_name)` : Object
- `get_boolean(column_idx)` : boolean
- `get_boolean(column_name)` : boolean
- `get_date(column_idx)` : `datetime.date`
- `get_date(column_name)` : `datetime.date`
- `get_float(column_idx)` : `Float`
- `get_float(column_name)` : `Float`
- `get_integer(column_idx)` : `Integer`
- `get_integer(column_name)` : `Integer`
- `get_list(column_idx)` : `List`
- `get_list(column_name)` : `List`
- `get_string(column_idx)` : `String`
- `get_string(column_name)` : `String`

- `get_time(column_idx) : datetime.time`
- `get_time(column_name) : datetime.time`
- `get_time_with_timezone(column_idx) : datetime.time`
- `get_time_with_timezone(column_name) : datetime.time`
- `get_timestamp(column_idx) : datetime.datetime`
- `get_timestamp(column_name) : datetime.datetime`
- `get_timestamp_with_timezone(column_idx) : datetime.datetime`
- `get_timestamp_with_timezone(column_name) : datetime.datetime`
- `get_value_type(column_idx) : Integer`
- `get_value_type(column_name) : Integer`
- `get_vertex_labels(column_idx) : List`
- `get_vertex_labels(column_name) : List`

See [Retrieving PGQL-on-RDBMS results](#) documentation for more information.

The following code samples illustrate cursor operations for iterating a result set using few of the cursor position and getter methods. These examples reference the query result set obtained in the [example](#) in the previous section.

```
Call first() and retrieve value for "FROM_ACCT_ID"
>>> pgql_result_set.first()
True
>>> pgql_result_set.get("FROM_ACCT_ID")
781.0

Call next() and retrieve value for "FROM_ACCT_ID"
>>> pgql_result_set.next()
True
>>> pgql_result_set.get("FROM_ACCT_ID")
978.0

Call last() and retrieve value for "FROM_ACCT_ID"
>>> pgql_result_set.last()
True
>>> pgql_result_set.get("FROM_ACCT_ID")
842.0

Call previous() and retrieve value for "FROM_ACCT_ID"
>>> pgql_result_set.previous()
True
>>> pgql_result_set.get("FROM_ACCT_ID")
838.0

Reset the result set and offset by 6. Then retrieve value for
"FROM_ACCT_ID"
>>> pgql_result_set.before_first()
>>> pgql_result_set.relative(6)
True
>>> pgql_result_set.get("FROM_ACCT_ID")
```



```
925.0

Reach the end of the result set and offset by -2. Then retrieve value for
"FROM_ACCT_ID"
>>> pgql_result_set.after_last()
>>> pgql_result_set.relative(-2)
True
>>> pgql_result_set.get("FROM_ACCT_ID")
838.0

Call absolute() and provide an absolute row value. Then retrieve value for
"FROM_ACCT_ID"
>>> pgql_result_set.absolute(3)
True
>>> pgql_result_set.get_float("FROM_ACCT_ID")
900.0

Get a specific row or a set of rows
>>> pgql_result_set.get_slice(0,1)
[781.0, 712.0, 1000.0]

>>> pgql_result_set.get_row(0)
[781.0, 712.0, 1000.0]
```

### Iterating a Result Set Using the Python Index Operator

You can also iterate through the query result set using the Python index operator as shown:

```
Retrieving a value from a tuple
>>> pgql_result_set[4, "double", "FROM_ACCT_ID"]
907.0

Retrieving a value using index value
>>> pgql_result_set[4].get("FROM_ACCT_ID")
907.0

Fetch a row or a set of rows
>>> pgql_result_set[0:1]
[781.0, 712.0, 1000.0]
```

### Iteration a Result Set Using a Python loop

Optionally, you can also iterate through the query result set using a Python loop. For example:

```
Using the result set as an iterator of a for loop
>>> for result in pgql_result_set:
 print(result)
[781.0, 712.0, 1000.0]
[190.0, 555.0, 1000.0]
[191.0, 329.0, 1000.0]
[198.8, 57.0, 1000.0]
[220.0, 441.0, 1000.0]
```

```
[251.0, 387.0, 1000.0]
[254.0, 188.0, 1000.0]
[259.0, 305.0, 1000.0]
[261.0, 145.0, 1000.0]
[263.0, 40.0, 1000.0]
```

### 6.9.1.6 Performance Considerations for PGQL Queries

Many factors affect the performance of PGQL queries in Oracle Database. The following are some recommended practices for query performance.

- [Query Optimizer Statistics](#)
- [Parallel Query Execution](#)
- [Optimizer Dynamic Sampling](#)
- [Bind Variables](#)
- [Path Queries](#)

#### Query Optimizer Statistics

Good, up-to-date query optimizer statistics are critical for query performance. Ensure that you run [OPG\\_APIS.ANALYZE\\_PG](#) after any significant updates to your property graph data.

#### Parallel Query Execution

Use parallel query execution to take advantage of Oracle's parallel SQL engine. Parallel execution often gives a significant speedup versus serial execution. Parallel execution is especially critical for path queries evaluated using the recursive WITH strategy.

See also the *Oracle Database VLDB and Partitioning Guide* for more information about parallel query execution.

#### Optimizer Dynamic Sampling

Due to the inherent flexibility of the graph data model, static information may not always produce optimal query plans. In such cases, dynamic sampling can be used by the query optimizer to sample data at run time for better query plans. The amount of data sampled is controlled by the dynamic sampling level used. Dynamic sampling levels range from 0 to 11. The best level to use depends on a particular dataset and workload, but levels of 2 (default), 6, or 11 often give good results.

See also Supplemental Dynamic Statistics in the *Oracle Database SQL Tuning Guide*.

#### Bind Variables

Use bind variables for constants whenever possible. The use of bind variables gives a very large reduction in query compilation time, which dramatically increases throughput for query workloads with queries that differ only in the constant values used. In addition, queries with bind variables are less vulnerable to injection attacks.

#### Path Queries

Path queries in PGQL that use the + (plus sign) or \* (asterisk) operator to search for arbitrary length paths require special consideration because of their high


computational complexity. You should use parallel execution and use the DISTINCT option for Recursive WITH (USE\_DIST\_RW=T) for the best performance. Also, for large, highly connected graphs, it is a good idea to use MAX\_PATH\_LEN=*n* to limit the number of repetitions of the recursive step to a reasonable number. A good strategy can be to start with a small repetition limit, and iteratively increase the limit to find more and more results.

## 6.9.2 Executing PGQL Queries Against Property Graph Views

This topic explains how you can execute PGQL queries directly against the property graph views on Oracle Database tables.

The PGQL query execution flow is shown in the following figure.

**Figure 6-2 PGQL on Property Graph Views in Oracle Database**



The basic execution flow is the same as in [PGQL on property graph schema tables](#). The only exception is that the PGQL query is translated into SQL statements using the internal metadata tables for property graph views.

- [PGQL Features Supported in Property Graph Views](#)
- [Performance Considerations for PGQL Queries](#)
- [Creating a Property Graph View](#)
- [Executing PGQL SELECT Queries](#)
- [Dropping A Property Graph View](#)

## 6.9.2.1 PGQL Features Supported in Property Graph Views

The following PGQL features are supported in property graph views:

- Recursive queries are supported for the following variable-length path finding goals:

- Reachability
- ANY
- ANY SHORTEST
- TOP k SHORTEST

- Recursive queries are supported for the following horizontal aggregations:

- LISTAGG

```
SELECT LISTAGG(src.first_name || ' ' || src.last_name, ',')
FROM MATCH TOP 2 SHORTEST ((n:Person) ((src)-[e:knows]->)*
(m:Person))
WHERE n.id = 1234
```

- SUM

```
SELECT SUM(e.weight + 3)
FROM MATCH TOP 2 SHORTEST ((n:Person) -[e:knows]->* (m:Person))
WHERE n.id = 1234
```

- COUNT

```
SELECT COUNT(e)
FROM MATCH TOP 2 SHORTEST ((n:Person) -[e:knows]->* (m:Person))
WHERE n.id = 1234
```

- AVG

```
SELECT AVG(dst.age)
FROM MATCH TOP 2 SHORTEST ((n:Person) (-[e:knows]->(dst))*
(m:Person))
WHERE n.id = 1234
```

- MIN (Only for property value or CAST expressions)

```
SELECT MIN(CAST(dst.age + 5 AS INTEGER))
FROM MATCH TOP 2 SHORTEST ((n:Person) (-[e:knows]->(dst))*
(m:Person))
WHERE n.id = 1234
```

- MAX (Only for property value or CAST expressions)

```
SELECT MAX(dst.birthdate)
FROM MATCH TOP 2 SHORTEST ((n:Person) (-[e:knows]->(dst))*
```

```
(m:Person))
WHERE n.id = 1234
```

- Both built-in Oracle Database functions and user defined functions (UDFs) are supported. For example:

- Assuming a table has a JSON column with values such as, {"name":"John", "age":43}:

```
SELECT JSON_VALUE(p.attributes, '$.name') AS name
FROM MATCH (p:Person)
WHERE JSON_VALUE(p.attributes, '$.age') > 35
```

- Assuming an Oracle Text index exists on a text column in a table:

```
SELECT n.text
FROM MATCH (n)
WHERE CONTAINS(n.text, 'cat', 1) > 0
```

- Assuming a UDF `updated_id` is registered with the graph server (PGX):

```
SELECT my.updated_id(n.ID) FROM MATCH(n) LIMIT 10
```

The following are a few PGQL features which are not supported:

- PGQL UPDATE queries are not supported for property graph views.
- The following PGQL SELECT features are not supported:

- The only quantifier supported for recursive queries is `*`.

If you attempt to use a different quantifier, it will result in an error as shown:

```
opg4j> String s = "SELECT id(a) FROM MATCH ANY SHORTEST ((a) -[e]->+
(b))";
s ==> "SELECT id(a) FROM MATCH ANY SHORTEST ((a) -[e]->+ (b))"
```

```
opg4j> PgqlStatement stmt = pgqlConn.createStatement();
stmt ==> oracle.pg.rdbms.pgql.PgqlExecution@27b9d5b7
```

```
opg4j> stmt.execute(s);
| Exception java.lang.UnsupportedOperationException: Only zero (0)
or more path quantifier is supported
```

- Use of bind variables in path expressions.

If you attempt to use a bind variable, it will result in an error as shown:

```
opg4j> String s = "SELECT id(a) FROM MATCH ANY SHORTEST (a) -[e]->*
(b) WHERE id(a) = ?";
s ==> "SELECT id(a) FROM MATCH ANY SHORTEST (a) -[e]->* (b) WHERE
id(a) = ?"
```

```
opg4j> PgqlPreparedStatement ps = pgqlConn.prepareStatement(s);
ps ==> oracle.pg.rdbms.pgql.PgqlExecution@7806db3f
```

```
opg4j> ps.setString(1, "PERSON(3)");
```

```
opg4j> ps.executeQuery();
| Exception java.lang.UnsupportedOperationException: Use of
bind variables for path queries is not supported
```

- Using subqueries.
- `in_degree` and `out_degree` functions.

 **Note:**

- See [Supported PGQL Features and Limitations](#) for a complete list of supported and unsupported PGQL features for PGQL on RDBMS (PG\_VIEW option)
- See [Performance Considerations for PGQL Queries](#) for details on recommended practices to enhance query performance for recursive queries.

## 6.9.2.2 Performance Considerations for PGQL Queries

The following are some recommended practices for query performance.

### Recursive Queries

The following indexes are recommended in order to speed up execution of recursive queries:

- For underlying VERTEX tables of the recursive pattern, an index on the key column
- For underlying EDGE tables of the recursive pattern, an index on the source key column

 **Note:**

You can also create index on (source key, destination key).

For example, consider the following CREATE PROPERTY GRAPH statement:

```
CREATE PROPERTY GRAPH people
 VERTEX TABLES(
 person
 KEY (id)
 LABEL person
 PROPERTIES(name, age)
)
 EDGE TABLES(
 knows
 key (person1, person2)
 SOURCE KEY (person1) REFERENCES person
 DESTINATION KEY (person2) REFERENCES person
 NO PROPERTIES
```

```
)
OPTIONS (PG_VIEW)
```

And also consider the following query:

```
SELECT COUNT(*)
FROM MATCH ANY SHORTEST ((n:Person) -[e:knows]->* (m:Person))
WHERE n.id = 1234
```

In order to improve performance of the recursive part of the preceding query, the following indexes must exist:

- CREATE INDEX <INDEX\_NAME> ON PERSON(ID)
- CREATE INDEX <INDEX\_NAME> ON KNOWS(PERSON1) or  
CREATE INDEX <INDEX\_NAME> ON KNOWS(PERSON1, PERSON2)

### Composite Vertex Keys

For composite vertex keys, query execution can be optimized with the creation of function-based indexes on the key columns:

- For underlying VERTEX tables of the recursive pattern, a function-based index on the comma-separated concatenation of key columns
- For underlying EDGE tables of the recursive pattern, a function-based index on the comma-separated concatenation of source key columns

#### Note:

You can also create index on (source key columns, destination key columns).

For example, consider the following CREATE PROPERTY GRAPH statement:

```
CREATE PROPERTY GRAPH people
 VERTEX TABLES (
 person
 KEY (id1, id2)
 LABEL person
 PROPERTIES(name, age)
)
 EDGE TABLES (
 knows
 key (id)
 SOURCE KEY (id1person1, id2person1) REFERENCES person
 DESTINATION KEY (id1person2, id2person2) REFERENCES person
 NO PROPERTIES
)
 OPTIONS (PG_VIEW)
```

And also consider the following query:

```
SELECT COUNT(*)
FROM MATCH ANY SHORTEST ((n:Person) -[e:knows]->* (m:Person))
WHERE n.id = 1234
```

In order to improve performance of the recursive part of the preceding query, the following indexes must exist:

- `CREATE INDEX <INDEX_NAME> ON PERSON (ID1 || ',' || ID2)`
- `CREATE INDEX <INDEX_NAME> ON KNOWS (ID1PERSON1 || ',' || ID2PERSON1)` or `CREATE INDEX <INDEX_NAME> ON KNOWS (ID1PERSON1 || ',' || ID2PERSON1, ID1PERSON2 || ',' || ID2PERSON2)`

If some of the columns in a composite vertex key is a string column, the column needs to be comma-escaped in the function-based index creation.

For example, if column `ID1` in table `PERSON` of the preceding example is of type `VARCHAR2(10)`, you need to escape the comma for the column as follows:

```
replace(ID1, ',', '\,')
```

So, the indexes to improve performance will result as shown:

- `CREATE INDEX <INDEX_NAME> ON PERSON (replace(ID1, ',', '\,') || ',' || ID2)`
- `CREATE INDEX <INDEX_NAME> ON KNOWS (replace(ID1PERSON1, ',', '\,') || ',' || ID2PERSON1)`

### 6.9.2.3 Creating a Property Graph View

You can create a property graph view using the `CREATE PROPERTY GRAPH` statement.

#### Creating a Property Graph View Using JShell

To create a property graph view using JShell, execute the following steps as shown:

1. Launch the JShell to work with the database as shown:

```
./bin/opg4j --no_connect
```

2. Connect to the Oracle Database as shown:

```
opg4j> var jdbcUrl="jdbc:oracle:thin:@<host_name>:<port>/<service>"
jdbcUrl ==> "jdbc:oracle:thin:@<host_name>:<port>/<db_service>"
opg4j> var conn =
DriverManager.getConnection(jdbcUrl,"<username>","<password>");
conn ==> oracle.jdbc.driver.T4CConnection@3e9a20e
```

In the preceding code:

- `<host_name>`: database host
- `<port>`: database port



- <service>: **database SID**
- <username>: **user name**
- <password>: **database password**

### 3. Create a PGQL Connection as shown:

```
opg4j> var pgqlConn = PgqlConnection.getConnection(conn)
pgqlConn ==> oracle.pg.rdbms.pgql.PgqlConnection@4301fa39
```

### 4. Execute the following comands as shown to create a property graph view:

```
opg4j> var pgqlStmt = pgqlConn.createStatement(); //create a PGQL
Statement
pgqlStmt ==> oracle.jdbc.driver.OracleStatementWrapper@6f976c
opg4j> String pgql =
...> "CREATE PROPERTY GRAPH <pgview> "
...> + "VERTEX TABLES (bank_nodes AS Accounts "
...> + "KEY (id) "
...> + "LABEL Accounts "
...> + "PROPERTIES (id, label) "
...> + ") "
...> + "EDGE TABLES (bank_edges_amt AS Transfers "
...> + "KEY (src_id, dest_id, amount) "
...> + "SOURCE KEY (src_id) REFERENCES Accounts "
...> + "DESTINATION KEY (dest_id) REFERENCES Accounts "
...> + "LABEL Transfers "
...> + "PROPERTIES (src_id, dest_id, amount, label) "
...> + ") OPTIONS (PG_VIEW) ";
opg4j> pgqlStmt.execute(pgql);
$8 ==> false
```

The property graph view is created.

## Creating a Property Graph View Using Java

The following example shows how to create a property graph view from the relational database tables using Java. The examples assumes the tables `bank_nodes` and `bank_edges_amt` already exist in the database.

```
import java.sql.Connection;
import java.sql.Statement;
import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlStatement;
import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;

/**
 * This example shows how to create a Property Graph View from relational
 * data stored in Oracle Database by executing a PGQL create statement.
 */
public class CreatePgView
{

 public static void main(String[] args) throws Exception
```

```

{
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String pgview = args[idx++];

 Connection conn = null;
 PgqlStatement pgqlStmt = null;

 try {

 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port + "/" + sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);

 // Create a PgqlStatement
 pgqlStmt = pgqlConn.createStatement();

 // Execute PGQL to create property graph view
 String pgql =
 "CREATE PROPERTY GRAPH " + pgview + " " +
 "VERTEX TABLES (bank_nodes as Accounts " +
 "KEY (id) " +
 "LABEL \"Accounts\"" +
 "PROPERTIES (id, label)" +
 ") " +
 "EDGE TABLES (bank_edges_amt as Transfers " +
 "KEY (src_id, dest_id, amount) " +
 "SOURCE KEY (src_id) REFERENCES Accounts " +
 "DESTINATION KEY (dest_id) REFERENCES Accounts " +
 "LABEL \"Transfers\"" +
 "PROPERTIES (src_id, dest_id, amount, label)" +
 ") OPTIONS (PG_VIEW) ";

 pgqlStmt.execute(pgql);

 }
 finally {
 // close the statement
 if (pgqlStmt != null) {
 pgqlStmt.close();
 }
 // close the connection
 }
}

```

```

 if (conn != null) {
 conn.close();
 }
 }
}
}
}

```

You can verify the property graph view creation by checking the metadata tables that get created in the Oracle Database.

## 6.9.2.4 Executing PGQL SELECT Queries

You can execute PGQL `SELECT` queries using the Java API in the `oracle.pg.rdbms.pgql` package. The Java interfaces `PgqlStatement` and `PgqlResultSet` are used in query execution.

### SELECT Query Execution on a Property Graph View Using JShell

1. Launch the JShell to work with the database as shown:

```
./bin/opg4j --no_connect
```

2. Execute the following commands to connect to the database and to run any `SELECT` query on a property graph view:

```

opg4j> var jdbcUrl="jdbc:oracle:thin:@<host_name>:<port>/<db_service>"
opg4j> var conn =
DriverManager.getConnection(jdbcUrl,"<username>","<password>");
opg4j> var pgqlConn = PgqlConnection.getConnection(conn)
opg4j> var pgqlStmt = pgqlConn.createStatement() //create a PGQL Statement
opg4j> String s = "SELECT n.id FROM MATCH (n:Accounts) ON <pgview> LIMIT
3"
opg4j> var rs = pgqlStmt.executeQuery(s)
opg4j> rs.print() //Prints the query result set
+-----+
| ID |
+-----+
| 434 |
| 435 |
| 436 |
+-----+

```

### SELECT Query Execution on a Property Graph View Using Java

The following example shows how to execute a `SELECT` query on a property graph view.

```

import java.sql.Connection;
import java.sql.Statement;
import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlResultSet;
import oracle.pg.rdbms.pgql.PgqlStatement;
import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;
/**

```

```

* This example shows how to execute a SELECT query on a property graph
view.
*/
public class ExecuteQueryOnPgView
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String pgview = args[idx++];

 Connection conn = null;
 PgqlStatement pgqlStmt = null;
 PgqlResultSet rs = null;

 try {
 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@"+host+": "+port + "/" + sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);

 // Create a PGQL Statement
 pgqlStmt = pgqlConn.createStatement();

 // Execute PGQL Query
 String query = "SELECT a,b,c,e1,e2,e3 FROM MATCH (a)-[e1]->(b)-
[e2]->(c)-[e3]->(a) ON " + pgview + " LIMIT 3";
 rs = pgqlStmt.executeQuery(query);

 // Print the results
 rs.print();
 }
 finally {
 // close the result set
 if (rs != null) {
 rs.close();
 }
 // close the statement
 if (pgqlStmt != null) {
 pgqlStmt.close();
 }
 }
 }
}

```

```

 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
}
}
}

```

The resulting output is as shown:

```

+-----+
+-----+
| A | B | C | E1 |
E2 | E3 | | |
+-----+
+-----+
| ACCOUNTS (998) | ACCOUNTS (781) | ACCOUNTS (71) | TRANSFERS (998,781,1000) |
TRANSFERS (781,71,1000) | TRANSFERS (71,998,1000) |
| ACCOUNTS (359) | ACCOUNTS (579) | ACCOUNTS (76) | TRANSFERS (359,579,1000) |
TRANSFERS (579,76,1000) | TRANSFERS (76,359,1000) |
| ACCOUNTS (6) | ACCOUNTS (580) | ACCOUNTS (82) | TRANSFERS (6,580,1000) |
TRANSFERS (580,82,1000) | TRANSFERS (82,6,1000) |
+-----+
+-----+

```

### 6.9.2.5 Dropping A Property Graph View

You can use PGQL to drop property graph views. When a `DROP PROPERTY GRAPH` statement is called, all the metadata tables for the property graph view are dropped.

#### Drop a Property Graph View Using JShell

1. Launch the JShell to work with the database as shown:

```
./bin/opg4j --no_connect
```

2. Execute the following commands to drop a property graph view:

```

opg4j> var jdbcUrl="jdbc:oracle:thin:@<host_name>:<port>/<db_service>"
opg4j> var conn =
DriverManager.getConnection(jdbcUrl,"<username>","<password>")
opg4j> var pgqlConn = PgqlConnection.getConnection(conn)
opg4j> var pgqlStmt = pgqlConn.createStatement() //create a PGQL Statement
opg4j> pgqlStmt.execute("DROP PROPERTY GRAPH <pgview>")
$9 ==> false

```

#### Drop a Property Graph View Using Java

The following example shows how to drop a property graph view.

```

import java.sql.Connection;
import java.sql.Statement;

```

```

import oracle.pg.rdbms.pgql.PgqlConnection;
import oracle.pg.rdbms.pgql.PgqlStatement;
import oracle.ucp.jdbc.PoolDataSourceFactory;
import oracle.ucp.jdbc.PoolDataSource;
/**
 * This example shows how to drop a property graph view.
 */
public class DropPgView
{

 public static void main(String[] args) throws Exception
 {
 int idx=0;
 String host = args[idx++];
 String port = args[idx++];
 String sid = args[idx++];
 String user = args[idx++];
 String password = args[idx++];
 String pgview = args[idx++];

 Connection conn = null;
 PgqlStatement pgqlStmt = null;

 try {
 //Get a jdbc connection
 PoolDataSource pds = PoolDataSourceFactory.getPoolDataSource();

 pds.setConnectionFactoryClassName("oracle.jdbc.pool.OracleDataSource");
 pds.setURL("jdbc:oracle:thin:@" + host + ":" + port + "/" + sid);
 pds.setUser(user);
 pds.setPassword(password);
 conn = pds.getConnection();
 conn.setAutoCommit(false);

 // Get a PGQL connection
 PgqlConnection pgqlConn = PgqlConnection.getConnection(conn);

 // Create PGQL Statement
 pgqlStmt = pgqlConn.createStatement();

 String query = "DROP PROPERTY GRAPH " + pgview;
 pgqlStmt.execute(query);

 }
 finally {
 // close the statement
 if (pgqlStmt != null) {
 pgqlStmt.close();
 }
 // close the connection
 if (conn != null) {
 conn.close();
 }
 }
 }
}

```

```
}
}
```

# 7

## Graph Visualization Application

The Graph Visualization application enables interactive exploration and visualization of property graphs. It can also visualize graphs stored in the database.

- [About the Graph Visualization Application](#)  
The Graph Visualization application is a single-page web application that works with the in-memory graph server (PGX).
- [How does the Graph Visualization Application Work](#)  
The Graph Visualization application exposes its own web interface and REST endpoint and can execute PGQL queries against the in-memory graph server (PGX) or the Oracle Database (PGQL on RDBMS).
- [Using the Graph Visualization Application](#)  
The principal points of entry for the Graph Visualization application are the query editor and the graph lists.
- [REST Endpoints for the Graph Visualization Application](#)
- [Kerberos Enabled Authentication for the Graph Visualization Application](#)  
The Graph Visualization application can authenticate users with Kerberos authentication enabled.

### 7.1 About the Graph Visualization Application

The Graph Visualization application is a single-page web application that works with the in-memory graph server (PGX).

The in-memory graph analytics server can be deployed in embedded mode or in Apache Tomcat or Oracle Weblogic Server. Graph Visualization application takes PGQL queries as an input and renders the result visually. A rich set of client-side exploration and visualization features can reveal new insights into your graph data.

Graph Visualization application works with the in-memory analytics server. It can visualize graphs that have been loaded into the in-memory analytics server, either preloaded when the in-memory analytics server is started, or loaded at run-time by a client application and made available through the `graph.publish()` API.

### 7.2 How does the Graph Visualization Application Work

The Graph Visualization application exposes its own web interface and REST endpoint and can execute PGQL queries against the in-memory graph server (PGX) or the Oracle Database (PGQL on RDBMS).

By default, it uses PGX and therefore requires a running PGX server to function. Alternatively, you can configure Graph Visualization application to directly talk to the database via PGQL on RDBMS. Graph Visualization application does not have any UI to create graphs, it can only visualize graphs which are already loaded into PGX or Oracle Database. See [REST Endpoints for the Graph Visualization Application](#) for more information on the graph visualization REST endpoints.



See [Enabling the Graph Visualization Application](#) for more information on starting the Graph Visualization application.


## 7.3 Using the Graph Visualization Application

The principal points of entry for the Graph Visualization application are the query editor and the graph lists.

When you start the graph visualization application, the graph list will be populated with the graphs loaded in PGX. To run queries against a graph, select that graph. The query lets you write PGQL queries that can be visualized. (PGQL is the SQL-like query language supported by the Graph Visualization application.)

Once the query is ready and the desired graph is selected, click the **Run** icon to execute the query. The following figure shows a query visualization identifying all edges that are directed edges from any vertex in the graph to any other vertex.

**Figure 7-1** Query Visualization



When a query is successful, the graph visualization is displayed, including nodes and their connections. You can right-click a node or connection to display tooltip information, and you can drag the nodes around.

- [Graph Visualization Modes](#)  
The buttons on the right let you switch between two modes: Graph Manipulation and Zoom/Move.
- [Graph Visualization Settings](#)  
You can click the **Settings** gear icon to display the Graph Visualization settings window.
- [Using the Geographical Layout](#)  
The Graph Visualization application offers a choice of layouts for rendering graphs. One of them is the Geographical layout that will show the graph (vertices and edges) on a global map.

- [Using Live Search](#)  
Live Search lets you to search the displayed graph and add live fuzzy search score to each item, so you can create a Highlight which visually shows the results of the search in the graph immediately.
- [Using URL Parameters to Control the Graph Visualization Application](#)  
You can provide the Graph Visualization application input data through URL parameters instead of using the form fields of the user interface.

## 7.3.1 Graph Visualization Modes

The buttons on the right let you switch between two modes: Graph Manipulation and Zoom/Move.

- **Graph Manipulation** mode lets you execute actions that modify the visualization. These actions include:
  - **Drop** removes selected vertices from visualization. Can also be executed from the tooltip.
  - **Group** selects multiple vertices and collapses them into a single one.
  - **Ungroup** selects a group of collapsed vertices and ungroups them.
  - **Expand** retrieves a configurable number of neighbors (hops) of selected vertices. Can also be executed from the tooltip.
  - **Focus**, like Expand, retrieves a configurable number of neighbors, but also drops all other vertices. Can also be executed from the tooltip.
  - **Undo** undoes the last action.
  - **Redo** redoes the last action.
  - **Reset** resets the visualization to the original state after the query.
- **Zoom/Move** mode lets you zoom in and out, as well as to move to another part of the visualization. The **Pan to Center** button resets the zoom and returns the view to the original one.


An additional mode, called **Sticky** mode, lets you cancel the action of dragging the nodes around.

## 7.3.2 Graph Visualization Settings

You can click the **Settings** gear icon to display the Graph Visualization settings window.

The settings window lets you modify some parameters for the visualization, and it has tabs for General, Visualization, and Highlights. The following figure shows this window, with the Visualization tab selected.

Figure 7-2 Graph Visualization Settings Window



The screenshot shows the 'Graph Visualization Settings Window' with three tabs: 'General', 'Visualization', and 'Highlights'. The 'Visualization' tab is selected. The settings are organized into two sections: 'General' and 'Layouts'.

**General Section:**

- Theme:** Light (selected) and Dark.
- Edge Style:** Straight and Curved (selected).
- Edge Marker:** Arrow (selected) and None.
- Similar Edges:** Collect and Keep (selected).
- Page Size:** 100, with up and down arrow buttons.

**Layouts Section:**

- Layout:** Force (selected in a dropdown menu).
- Edge Distance:** Slider set to 120.
- Force Strength:** Slider set to -30.
- Velocity Decay:** Slider set to 0.3.
- Vertex Padding:** Slider set to 40.

The **General tab** includes the following:

- **Number of hops:** The configurable number of hops for the expand and focus actions.
- **Truncate label:** Truncates the label if it exceeds the maximum length.
- **Max. visible label length:** Maximum length before truncating.
- **Show Label On Hover:** Controls whether the label is shown on hover.
- **Display the graph legend:** Controls whether the legend is displayed.

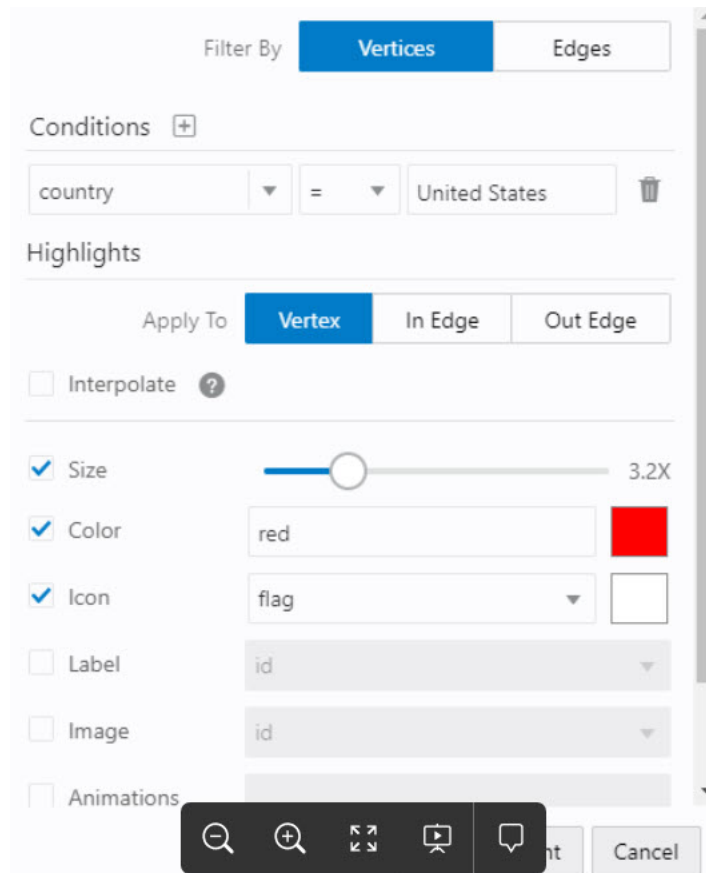
The **Visualization tab** includes the following:

- **Theme:** Select a light or dark mode.
- **Edge Style:** Select straight or curved edges.
- **Edge Marker:** Select arrows or no edge marker. This only applies to directed edges.
- **Similar Edges:** Select keep or collect.
- **Page Size:** Specify how many vertices and edges are displayed per page.
- **Layouts:** Select between different layouts (random, grid, circle, concentric, ...).
- **Vertex Label:** Select which property to use as the vertex label.
- **Vertex Label Orientation:** Select the relative position of the vertex label.

- **Edge Label:** Select which property to use as the edge label.

The **Highlights tab** includes customization options that let you modify the appearance of edges and vertices. Highlighting can be applied based on conditions (filters) on single or multiple elements. The following figure shows a condition (`country = United States`) and visual highlight options for vertices.

**Figure 7-3 Highlights Options for Vertices**



A filter for highlights can contain multiple conditions on any property of the element. The following conditions are supported.

- = (equal to)
- < (less than)
- <= (less than or equal to)
- > (greater than)
- >= (greater than or equal to)
- != (not equal to)
- ~ (filter is a regular expression)
- \* (any: like a wildcard, can match to anything)

The visual highlight customization options include:

- Edges:
  - Width
  - Color
  - Label
  - Style
  - Animations
- Vertices:
  - Size
  - Color
  - Icon
  - Label
  - Image
  - Animations

You can export and import highlight options by clicking the Save and Import buttons in the main window. **Save** lets you persist the highlight options, and **Load** lets you apply previously saved highlight options.

When you click **Save**, a file is saved containing a JSON object with the highlights configuration. Later, you can load that file to restore the highlights of the saved session.

### 7.3.3 Using the Geographical Layout

The Graph Visualization application offers a choice of layouts for rendering graphs. One of them is the Geographical layout that will show the graph (vertices and edges) on a global map.

The following figure shows a graph rendered on a geographical layout in the Graph Visualization application:

**Figure 7-4 Geographical Layout**



In order to view your vertices on a map, they must include a geographical location, in the form of a pair of properties that contain the longitude and latitude coordinates for that vertex. For example:


| iata | city        | longitude    | latitude    |
|------|-------------|--------------|-------------|
| SIN  | Singapore   | 103.994003   | 1.35019     |
| LAX  | Los Angeles | -118.4079971 | 33.94250107 |
| MUC  | Munich      | 11.7861      | 48.353802   |
| CDG  | Paris       | 2.55         | 49.012798   |
| LHR  | London      | -0.461941    | 51.4706     |

**Note:**

You can use any name for the longitude and latitude properties (such as X and Y, or long and lat). But, you must ensure that the longitude/latitude pair are in the WGS84 system (GPS coordinates), and the coordinates are expressed in decimal degrees.

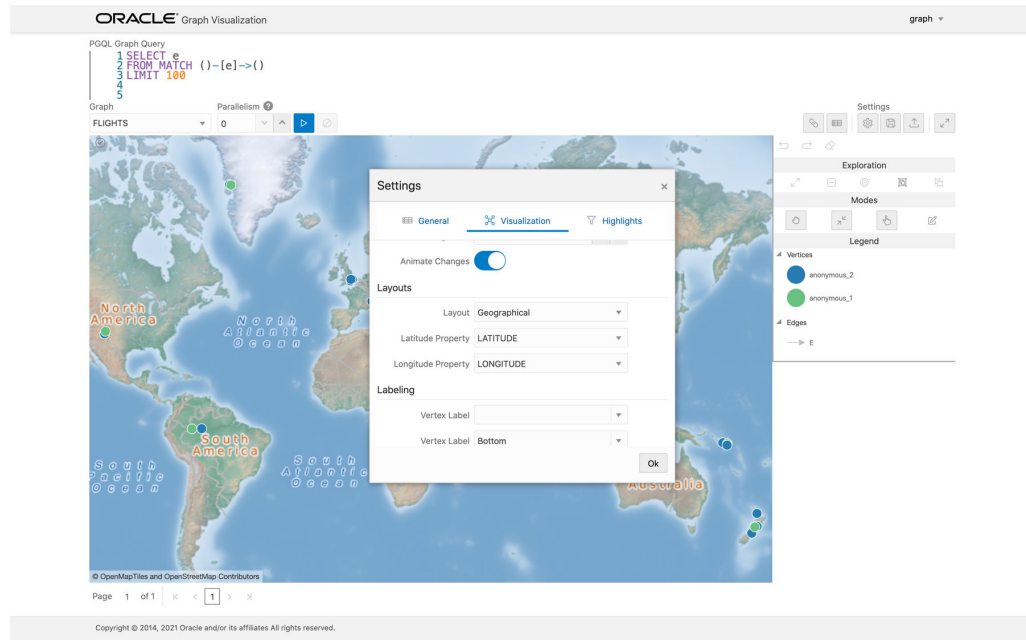
You can select the geographical layout in the Graph Visualization settings window as shown:

**Figure 7-5 Setting Geographical Layout**



Then, select the properties in your vertices that contain the geographical coordinates as shown:

Figure 7-6 Selecting the Coordinates for the Geographical layout



You can now move around the map and zoom in/out using your mouse or trackpad. From now on, whenever you enter a new PGQL query, the map will automatically center and zoom the vertices returned by the query.

### 7.3.4 Using Live Search

Live Search lets you to search the displayed graph and add live fuzzy search score to each item, so you can create a Highlight which visually shows the results of the search in the graph immediately.

If you run a query, and a graph is displayed, you can add the live search, which is on the settings dialog. On the bottom of the General tab, you will see these options.

- **Enable Live Search:** Enables the Live Search feature, adds the search input to the visualization, and lets you further customize the search.
- **Enable Search In:** You can select whether you want to search the properties of Vertices, Edges, or both.
- **Properties To Search:** Based on what you selected for Enable Search In, you can set one or more properties to search in. For example, if you disable the search for edges but you had a property from edges selected, it will be stored and added back when you enable search for the edges again. (This also works for vertices.)
- **Advanced Settings:** You can fine-tune the search even more. Each of the advanced options is documented with context help, visible upon enabling.
  - **Location:** Determines approximately where in the text the pattern is expected to be found.
  - **Distance:** Determines how close the match must be to the fuzzy location (specified by location). An exact letter match which is distance characters away from the fuzzy location would score as a complete mismatch. A distance

of 0 requires the match be at the exact location specified, a distance of 1000 would require a perfect match to be within 800 characters of the location to be found using a threshold of 0.8.

- **Maximum Pattern Length:** The maximum length of the pattern. The longer the pattern (that is, the search query), the more intensive the search operation will be. Whenever the pattern exceeds this value, an error will be thrown.
- **Min Char Match:** The minimum length of the pattern. Whenever the pattern length is below this value, an error will be thrown.

When the search is enabled, the input will be displayed in the top left part of the Graph Visualization component. If you start typing, the search will add a score to every vertex or edge, based on the settings and the search match.

To be able to see the results visually, you have to add a **Highlight** with interpolation set to a **Live Search** score and other settings based on the desired visual change.

### 7.3.5 Using URL Parameters to Control the Graph Visualization Application

You can provide the Graph Visualization application input data through URL parameters instead of using the form fields of the user interface.

If you supply the parameters in the URL, the Graph Visualization application automatically executes the specified query and hides the input form fields from the screen, so only the resulting visualization output is visible. This feature is useful if you want to embed the resulting graph visualization into an existing application, such as through an iframe.

The following table specifies the available URL parameters:

**Table 7-1 Available URL Parameters**

| Parameter Name | Value (must be URL encoded)   | Type   | Optional?                                         |
|----------------|-------------------------------|--------|---------------------------------------------------|
| graph          | Graph name                    | string | No                                                |
| parallelism    | Degree of parallelism desired | number | Yes (defaults to server-side default parallelism) |
| query          | PQL query                     | string | No                                                |

The following URL shows an example of visualizing the PGQL query `SELECT v, e MATCH (v) -[e]-> () LIMIT 10` on graph `myGraph` with parallelism 4:

```
https://myhost:7007/ui/?query=SELECT%20v%2C%20e%20MATCH%20%28v%29%20-%5Be%5D-%3E%20%28%29%20LIMIT%2010&graph=myGraph¶llelism=4
```

## 7.4 REST Endpoints for the Graph Visualization Application

This section explains all the REST endpoints through which you can perform various operations using the Graph Visualization Application.

The following are the available REST endpoints:



 **Note:**

The examples shown in the REST endpoints assume that:

- The PGX server is up and running on `https://localhost:7007`.
- Linux with `cURL` is installed. `cURL` is used to demonstrate how to access the `graph.publish` API using the CA certificate for verifying the graph server.

- [Login](#)
- [List Graphs](#)
- [Run a PGQL Query](#)
- [Get User](#)
- [Asynchronous REST Endpoints](#)

## 7.4.1 Login

**HTTP Request:** POST `https://localhost:7007/ui/v1/login/`

**Authentication:** Uses cookie-based authentication.

**Table 7-2 Parameters**

| Parameter                 | Parameter Type | Value                                                                                                                                                                                                                      | Required                                                                                                                         |
|---------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <code>Content-type</code> | Header         | <code>application/json</code>                                                                                                                                                                                              | Yes                                                                                                                              |
| <code>username</code>     | Body           | <code>&lt;username&gt;</code>                                                                                                                                                                                              | Yes                                                                                                                              |
| <code>password</code>     | Body           | <code>&lt;password&gt;</code>                                                                                                                                                                                              | Yes                                                                                                                              |
| <code>baseUrl</code>      | Body           | <code>&lt;baseUrl&gt;</code> to point to the graph server (PGX) or the database                                                                                                                                            | Optional. If empty, Graph Visualization will use the <code>pgx.base_url</code> parameter value in the <code>web.xml</code> file. |
| <code>pgqlDriver</code>   | Body           | Valid PGQL driver configuration values are: <ul style="list-style-type: none"> <li>• <code>pgxDriver</code>: for PGQL on the graph server (PGX)</li> <li>• <code>pgqlDriver</code>: for PGQL on Oracle Database</li> </ul> | Yes                                                                                                                              |
| <code>sessionId</code>    | Body           | <code>sessionId</code> from graph server (PGX)                                                                                                                                                                             | Optional                                                                                                                         |

### Request

The following `curl` command signs the user in to the Graph Visualization application:

```
curl --cacert /etc/oracle/graph/ca_certificate.pem -c cookie.txt -X POST -H
"Content-Type: application/json" -d '{"username": "<username>", "password":
"<password>", "pgqlDriver": "<pgqlDriver>", "baseUrl": "<baseUrl>",
"sessionId": "<sessionId>" }' https://localhost:7007/ui/v1/login/
```

**Response:** The username used for the login. For example:

```
"oracle"
```

On successful login, the server session cookie is stored in a cookie file, `cookie.txt`. Use this cookie file, in the subsequent calls to the API.

## 7.4.2 List Graphs

**HTTP Request:** GET `https://localhost:7007/ui/v1/graphs`  
**Request**

The following `curl` command lists all the graphs that belong to the user:

```
curl --cacert /etc/oracle/graph/ca_certificate.pem -b cookie.txt 'https://
localhost:7007/ui/v1/graphs'
```

**Response:** The list of the available graphs for the current user. For example:

```
["hr", "bank_graph_analytics"]
```

## 7.4.3 Run a PGQL Query

**HTTP Request:** GET `https://localhost:7007/ui/v1/query?pgql=<PGQL
query>&graph=<graph>&parallelism=<value>&size=<size value>`

**Table 7-3 Query Parameters**

| Parameter                | Description                                           | Values                                 |
|--------------------------|-------------------------------------------------------|----------------------------------------|
| <code>pgql</code>        | PGQL query string                                     | <code>&lt;PGQL query&gt;</code>        |
| <code>graph</code>       | Name of the graph                                     | <code>&lt;graph_name&gt;</code>        |
| <code>parallelism</code> | Degree of Parallelism                                 | <code>&lt;parallelism_value&gt;</code> |
| <code>size</code>        | Fetch size (= the number of rows) of the query result | <code>&lt;size_value&gt;</code>        |

**Request**

The following `curl` command executes PGQL Query on a property graph:

```
curl --cacert /etc/oracle/graph/ca_certificate.pem -b cookie.txt
'https://localhost:7007/ui/v1/query?pgql=SELECT%20e%0AMATCH%20()-
%5Be%5D-%3E()%0ALIMIT%205&graph=hr¶llelism=&size=100'
```

**Response:** The PGQL query result in JSON format.

```
{
 "name": "bank_graph_analytics_2",
 "resultSetId": "pgql_14",
 "graph": {
 "idType": "number",
 "vertices": [
 {
 "_id": "1",
 "p": [],
 "l": [
 "Accounts"
],
 "g": [
 "anonymous_1"
]
 },
 {
 "_id": "418",
 "p": [],
 "l": [
 "Accounts"
],
 "g": [
 "anonymous_2"
]
 },
 {
 "_id": "259",
 "p": [],
 "l": [
 "Accounts"
],
 "g": [
 "anonymous_2"
]
 }
],
 "edges": [
 {
 "_id": "0",
 "p": [
 {
 "n": "AMOUNT",
 "v": "1000.0",
 "s": false
 }
]
 }
]
 }
}
```

```

],
 "l": [
 "Transfers"
],
 "g": [
 "e"
],
 "s": "1",
 "d": "259",
 "u": false
 },
 {
 "_id": "1",
 "p": [
 {
 "n": "AMOUNT",
 "v": "1000.0",
 "s": false
 }
],
 "l": [
 "Transfers"
],
 "g": [
 "e"
],
 "s": "1",
 "d": "418",
 "u": false
 }
],
"paths": [],
"totalNumResults": 2
},
"table":
"e\nPgxE[provider=Transfers,ID=0]\nPgxE[provider=Transfers,ID=1]"
}

```

## 7.4.4 Get User

**HTTP Request:** GET <https://localhost:7007/ui/v1/user>

### Request

The following `curl` command gets the name of the current user:

```
curl --cacert /etc/oracle/graph/ca_certificate.pem -b cookie.txt 'https://localhost:7007/ui/v1/user'
```

**Response:** The name of the current user. For example:

```
"oracle"
```

## 7.4.5 Asynchronous REST Endpoints

Graph Visualization REST endpoints support cancellation of queries.

In order to be able to cancel queries, you need to send the query using the following asynchronous REST endpoints:

- [Run a PGQL Query Asynchronously](#)
- [Check a Query Completion](#)
- [Cancel a Query Execution](#)
- [Retrieve a Query Result](#)

### 7.4.5.1 Run a PGQL Query Asynchronously

**HTTP Request:** GET `https://localhost:7007/ui/v1/async-query?pgql=<PGQL query>&graph=<graph>&parallelism=<value>&size=<size value>`

See [Table 7-3](#) for more information on query parameters.

#### Request

The following `curl` command executes a PGQL query asynchronously on a property graph:

```
curl --cacert /etc/oracle/graph/ca_certificate.pem -b cookie.txt
'https://localhost:7007/ui/v1/async-query?pgql=SELECT%20e%0AMATCH%20()-%5Be%5D-%3E()%0ALIMIT%205&graph=hr¶llelism=&size=100'
```

**Response:** None.



#### Note:

An error message will be returned in case the query is malformed or if the graph does not exist.

### 7.4.5.2 Check a Query Completion

**HTTP Request:** GET `https://localhost:7007/ui/v1/async-query-complete`

#### Request

The following `curl` command checks if the PGQL query execution is completed:

```
curl --cacert /etc/oracle/graph/ca_certificate.pem -b cookie.txt
'https://localhost:7007/ui/v1/async-query-complete'
```

**Response:** A boolean that indicates if the query execution is completed. For example,

```
true
```

 **Note:**

You do not have to specify any request ID, as the currently executing query is attached to your HTTP session. You can only have one query executing per session. For concurrent query execution, create multiple HTTP sessions by logging in multiple times.

### 7.4.5.3 Cancel a Query Execution

**HTTP Request:** DELETE https://localhost:7007/ui/v1/async-query

#### Request

The following `curl` command cancels a currently executing PGQL Query on a property graph:

```
curl -X DELETE --cacert /etc/oracle/graph/ca_certificate.pem -b cookie.txt 'https://localhost:7007/ui/v1/async-query'
```

**Response:** Confirmation of the cancellation or an error message if the query has already completed execution.

### 7.4.5.4 Retrieve a Query Result

**HTTP Request:** GET https://localhost:7007/ui/v1/async-result

 **Note:**

The endpoint, GET https://localhost:7007/ui/v1/async-result?pgql=<PGQL query>&graph=<graph>&parallelism=<value>&size=<size value>, to retrieve a query result is deprecated:

```
curl --cacert /etc/oracle/graph/ca_certificate.pem -b cookie.txt 'https://localhost:7007/ui/v1/async-result?pgql=SELECT%20e%0AMATCH%20()-%5Be%5D-%3E()%0ALIMIT%205&graph=hr¶llelism=&size=100'
```

#### Request

The following `curl` command retrieves the result of a successfully completed query:

```
curl --cacert /etc/oracle/graph/ca_certificate.pem -b cookie.txt 'https://localhost:7007/ui/v1/async-result'
```

**Response:** The PGQL query result in JSON format.

```
{
 "name": "bank_graph_analytics_2",
 "resultSetId": "pgql_14",
 "graph": {
 "idType": "number",
 "vertices": [
 {
 "_id": "1",
 "p": [],
 "l": [
 "Accounts"
],
 "g": [
 "anonymous_1"
]
 },
 {
 "_id": "418",
 "p": [],
 "l": [
 "Accounts"
],
 "g": [
 "anonymous_2"
]
 },
 {
 "_id": "259",
 "p": [],
 "l": [
 "Accounts"
],
 "g": [
 "anonymous_2"
]
 }
],
 "edges": [
 {
 "_id": "0",
 "p": [
 {
 "n": "AMOUNT",
 "v": "1000.0",
 "s": false
 }
],
 "l": [
 "Transfers"
],
 "g": [
 "e"
]
 }
]
 }
}
```

```

 "s": "1",
 "d": "259",
 "u": false
 },
 {
 "_id": "1",
 "p": [
 {
 "n": "AMOUNT",
 "v": "1000.0",
 "s": false
 }
],
 "l": [
 "Transfers"
],
 "g": [
 "e"
],
 "s": "1",
 "d": "418",
 "u": false
 }
],
"paths": [],
"totalNumResults": 2
},
"table":
"e\nPgxEde[provider=Transfers,ID=0]\nPgxEdge[provider=Transfers,ID=1]"
}

```

## 7.5 Kerberos Enabled Authentication for the Graph Visualization Application

The Graph Visualization application can authenticate users with Kerberos authentication enabled.

Graph Visualization provides two different drivers to log in:

- **Graph Server (PGX) Driver:** To send your credentials (Kerberos ticket) to Graph Server.
- **Database Driver:** To send your credentials (Kerberos ticket) directly to the database.
- [Prerequisite Requirements for Kerberos Authentication](#)
- [Preparing the Graph Visualization Application for Kerberos Authentication](#)

### 7.5.1 Prerequisite Requirements for Kerberos Authentication

The system requirements for the respective PGQL drivers are as follows:

- **Graph Server (PGX) Driver:** See [Prerequisite Requirements](#) for enabling Kerberos authentication on the graph server (PGX).
- **Database Driver:**



- The database must have Kerberos authentication enabled. See Configuring Kerberos Authentication for more information.
- Both the database and the Kerberos Authentication Server need to be reachable from the host where the Graph Visualization application is running.
- The database must be prepared for graph server authentication. That is, relevant graph roles have been granted to users who will log into the Graph Visualization application.

## 7.5.2 Preparing the Graph Visualization Application for Kerberos Authentication

In order to use Kerberos authentication, you must enter your Active Directory credentials in the Graph Visualization application login page. To enable Kerberos authentication for the Graph Visualization application, follow the steps shown:

1. Locate the `web.xml` file for your installation.

You can locate the `WEB-INF/web.xml` inside the Graph Visualization WAR file for your installation as shown in the following table:

**Table 7-4 Location of `WEB-INF/web.xml` file**

| Type of Installation          | WAR file                                                                                                                                                   | Location                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Standalone installation (RPM) | <code>graphviz-&lt;version&gt;-pgviz&lt;graphviz-version&gt;.war</code>                                                                                    | <code>/opt/oracle/graph/graphviz</code>                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Apache Tomcat Deployment:     | <code>graphviz-&lt;version&gt;-pgviz&lt;graphviz-version&gt;-tomcat.war</code><br><version> denotes the downloaded Oracle Graph Server and Client version. | <ol style="list-style-type: none"> <li>a. Download <code>oracle-graph-webapps-&lt;version&gt;.zip</code> from <a href="#">Oracle Software Delivery Cloud</a></li> <li>b. Unzip the file into a directory of your choice.</li> <li>c. Locate the <code>.war</code> file for deploying the Graph Visualization application to Tomcat. It follows the naming pattern: <code>graphviz-&lt;version&gt;-pgviz&lt;graphviz-version&gt;-tomcat.war</code></li> </ol> |

**Table 7-4 (Cont.) Location of WEB-INF/web.xml file**

| Type of Installation              | WAR file                                                                                                                       | Location                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-----------------------------------|--------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oracle WebLogic Server Deployment | graphviz-<version>-pgviz<graphviz-version>-wls.war<br><version> denotes the downloaded Oracle Graph Server and Client version. | <ol style="list-style-type: none"> <li>a. Download <code>oracle-graph-webapps-&lt;version&gt;.zip</code> from <a href="#">Oracle Software Delivery Cloud</a></li> <li>b. Unzip the file into a directory of your choice.</li> <li>c. Locate the <code>.war</code> file for deploying the Graph Visualization application to Oracle WebLogic Server. It follows the naming pattern:<br/>graphviz-&lt;version&gt;-pgviz&lt;graphviz-version&gt;-wls.war</li> </ol> |

2. Extract the appropriate WAR file to a directory of your choice by executing the following command:

```
unzip graphviz-*.war -d <war-file-extraction-path>
```

3. Locate and open the WEB-INF/web.xml file for update using any file editor of your choice. For example:

```
cd <war-file-extraction-path>
vi WEB-INF/web.xml
```

4. Enable the `graphviz.driver.auth.kerberos` parameter as shown:

```
<context-param>
 <param-name>graphviz.driver.auth.kerberos</param-name>
 <param-value>true</param-value>
</context-param>
```

Setting this flag **true** initiates the Graph Visualization application to install its own `okinit` package.

5. Optionally, set the cache directory that will be used by the Graph Visualization application to temporarily store Kerberos tickets given by clients as shown

```
<context-param>
 <param-name>graphviz.driver.auth.kerberos.cache_dir</param-name>
 <param-value>/dev/shm/graph_cache</param-value>
</context-param>
```

The default value is `/dev/shm/graph_cache`. If the directory does not exist, it will be automatically created upon server startup.

- Optionally, set the maximum amount of concurrent Kerberos active sessions in the Graph Visualization application.

```
<context-param>
 <param-name>graphviz.driver.auth.kerberos.max_cache_size</
param-name>
 <param-value>64</param-value>
</context-param>
```

- Optionally, modify the directory where `okinit` package will be installed, by updating the following parameter:

```
<context-param>
 <param-name>graphviz.driver.auth.kerberos.okinit-directory</
param-name>
 <param-value>/tmp</param-value>
</context-param>
```

 **Note:**

The default value is `/tmp` and you must have executable permission for the directory.

- Optionally, set the following parameter if there is a location for an existing `okinit` package on your machine. In this case, the GraphVisualization application will not install its own `okinit` package.

```
<context-param>
 <param-
name>graphviz.driver.auth.kerberos.graphviz.driver.auth.okinit-
location</param-name>
 <param-value></param-value>
</context-param>
```

 **Note:**

The GraphVisualization application must have executable permission for the directory location.

- Finally, after all the preceding updates, repackage the `WAR` file by executing the following commands:

```
cd <war-file-extraction-path>
jar -cvf <war-file-name> *
```

- Redeploy the `WAR` file to the appropriate directory for your installation.

Kerberos authentication is enabled for the Graph Visualization Application.

# 8

## Using the Machine Learning Library (PgXML) for Graphs

The in-memory graph server (PGX) provides a machine learning library `oracle.pgx.api.mllib`, which supports graph-empowered machine learning algorithms.

The following machine learning algorithms are currently supported:

- [Using the DeepWalk Algorithm](#)  
**DeepWalk** is a widely employed vertex representation learning algorithm used in industry.
- [Using the Supervised GraphWise Algorithm](#)  
**Supervised GraphWise** is an inductive vertex representation learning algorithm which is able to leverage vertex feature information. It can be applied to a wide variety of tasks, including vertex classification and link prediction.
- [Using the Unsupervised GraphWise Algorithm](#)  
**Unsupervised GraphWise** is an unsupervised inductive vertex representation learning algorithm which is able to leverage vertex information. The learned embeddings can be used in various downstream tasks including vertex classification, vertex clustering and similar vertex search.
- [Using the Pg2vec Algorithm](#)  
**Pg2vec** learns representations of graphlets (partitions inside a graph) by employing edges as the principal learning units and thereby packing more information in each learning unit (as compared to employing vertices as learning units) for the representation learning task.

### 8.1 Using the DeepWalk Algorithm

**DeepWalk** is a widely employed vertex representation learning algorithm used in industry.

It consists of two main steps:

1. First, the random walk generation step computes random walks for each vertex (with a pre-defined walk length and a pre-defined number of walks per vertex).
2. Second, these generated walks are fed to a **Word2vec** algorithm to generate the vector representation for each vertex (which is the word in the input provided to the Word2vec algorithm). See [KDD paper](#) for more details on DeepWalk algorithm.

DeepWalk creates vertex embeddings for a specific graph and cannot be updated to incorporate modifications on the graph. Instead, a new DeepWalk model should be trained on this modified graph. Lastly, it is important to note that the memory consumption of the DeepWalk model is  $O(2n*d)$  where  $n$  is the number of vertices in the graph and  $d$  is the embedding length.

The following describes the usage of the main functionalities of DeepWalk in in-memory PGX using [DBpedia](#) graph as an example with 8,637,721 vertices and 165,049,964 edges:

- [Loading a Graph](#)

- [Building a Minimal DeepWalk Model](#)
- [Building a Customized DeepWalk Model](#)
- [Training a DeepWalk Model](#)
- [Getting the Loss Value For a DeepWalk Model](#)
- [Computing Similar Vertices for a Given Vertex](#)
- [Computing Similar Vertices for a Vertex Batch](#)
- [Getting All Trained Vertex Vectors](#)
- [Storing a Trained DeepWalk Model](#)
- [Loading a Pre-Trained DeepWalk Model](#)
- [Destroying a DeepWalk Model](#)

## 8.1.1 Loading a Graph

The following describes the steps for loading a graph:

1. Create a **Session** and an **Analyst**.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
cd /opt/oracle/graph/
./bin/opg4j
// starting the shell will create an implicit session and analyst
```

### Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.mllib.DeepWalkModel;
import oracle.pgx.api.frames.*;
...
PgxSession session = Pgx.createSession("my-session");
Analyst analyst = session.createAnalyst();
```

### Python

```
session = pypgx.get_session(session_name="my-session")
analyst = session.create_analyst()
```

- 
2. Load the **graph**.

 **Note:**

Though the DeepWalk algorithm implementation can be applied to directed or undirected graphs, currently only undirected random walks are considered.

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var graph = session.readGraphWithProperties("<path>/<graph.json>")
```

## Java

```
PgxGraph graph = session.readGraphWithProperties("<path>/<graph.json>");
```

## Python

```
graph = session.read_graph_with_properties("<path>/<graph.json>")
```

## 8.1.2 Building a Minimal DeepWalk Model

You can build a DeepWalk model using the minimal configuration and default hyperparameters as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var model = analyst.deepWalkModelBuilder().
 setWindowSize(3).
 setWalksPerVertex(6).
 setWalkLength(4).
 build()
```

## Java

```
DeepWalkModel model = analyst.deepWalkModelBuilder()
 .setWindowSize(3)
 .setWalksPerVertex(6)
 .setWalkLength(4)
 .build();
```

## Python

```
model =
analyst.deepwalk_builder(window_size=3,walks_per_vertex=6,walk_length=4
)
```

---

### 8.1.3 Building a Customized DeepWalk Model

You can build a DeepWalk model using customized hyper-parameters as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg4j> var model = analyst.deepWalkModelBuilder().
 setMinWordFrequency(1).
 setBatchSize(512).
 setNumEpochs(1).
 setLayerSize(100).
 setLearningRate(0.05).
 setMinLearningRate(0.0001).
 setWindowSize(3).
 setWalksPerVertex(6).
 setWalkLength(4).
 setSampleRate(0.00001).
 setNegativeSample(2).
 setValidationFraction(0.01).
 build()
```

#### Java

```
DeepWalkModel model= analyst.deepWalkModelBuilder()
 .setMinWordFrequency(1)
 .setBatchSize(512)
```

```
.setNumEpochs(1)
.setLayerSize(100)
.setLearningRate(0.05)
.setMinLearningRate(0.0001)
.setWindowSize(3)
.setWalksPerVertex(6)
.setWalkLength(4)
.setSampleRate(0.00001)
.setNegativeSample(2)
.setValidationFraction(0.01)
.build();
```

## Python

```
model = analyst.deepwalk_builder(min_word_frequency=1,
 batch_size=512,num_epochs=1,
 layer_size=100,
 learning_rate=0.05,
 min_learning_rate=0.0001,
 window_size=3,
 walks_per_vertex=6,
 walk_length=4,
 sample_rate=0.00001,
 negative_sample=2,
 validation_fraction=0.01)
```

---

See [DeepWalkModelBuilder](#) in Javadoc for more explanation for each builder operation along with the default values.

## 8.1.4 Training a DeepWalk Model

You can train a DeepWalk model with the specified default or customized settings as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
opg4j> model.fit(graph)
```

### Java

```
model.fit(graph);
```



## Python

```
model.fit(graph)
```

---

### 8.1.5 Getting the Loss Value For a DeepWalk Model

You can fetch the loss value on a specified fraction of training data, that is set in builder using `setValidationFraction` as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg4j> var loss = model.getLoss()
```

#### Java

```
double loss = model.getLoss();
```

#### Python

```
loss = model.loss
```

---

### 8.1.6 Computing Similar Vertices for a Given Vertex

You can fetch the  $k$  most similar vertices for a given vertex as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var similars = model.computeSimilars("Albert_Einstein", 10)
opg4j> similars.print()
```

## Java

```
PgxFrame similars = model.computeSimilars("Albert_Einstein", 10);
similars.print();
```

## Python

```
similars = model.compute_similars("Albert_Einstein",10)
similars.print()
```

---

Searching for similar vertices for [Albert\\_Einstein](#) using the trained model, will result in the following output:

```
+-----+
| dstVertex | similarity |
+-----+
| Albert_Einstein | 1.0000001192092896 |
| Physics | 0.8664291501045227 |
| Werner_Heisenberg | 0.8625140190124512 |
| Richard_Feynman | 0.8496938943862915 |
| List_of_physicists | 0.8415523767471313 |
| Physicist | 0.8384397625923157 |
| Max_Planck | 0.8370327353477478 |
| Niels_Bohr | 0.8340970873832703 |
| Quantum_mechanics | 0.8331197500228882 |
| Special_relativity | 0.8280861973762512 |
+-----+
```

### 8.1.7 Computing Similar Vertices for a Vertex Batch

You can fetch the  $k$  most similar vertices for a list of input vertices as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var vertices = new ArrayList()
opg4j> vertices.add("Machine_learning")
opg4j> vertices.add("Albert_Einstein")
opg4j> batchedSimilar = model.computeSimilar(vertices, 10)
opg4j> batchedSimilar.print()
```

## Java

```
List vertices = Arrays.asList("Machine_learning","Albert_Einstein");
PgxFrame batchedSimilar = model.computeSimilar(vertices,10);
batchedSimilar.print();
```

## Python

```
vertices = ["Machine_learning","Albert_Einstein"]
batched_similar = model.compute_similar(vertices,10)
batched_similar.print()
```

The following describes the output result:

```
+-----+
| srcVertex | dstVertex | similarity |
+-----+
| Machine_learning | Machine_learning | 1.0000001192092896 |
| Machine_learning | Data_mining | 0.9070799350738525 |
| Machine_learning | Computer_science | 0.8963605165481567 |
| Machine_learning | Unsupervised_learning | 0.8828719854354858 |
| Machine_learning | R_(programming_language) | 0.8821185827255249 |
| Machine_learning | Algorithm | 0.8819515705108643 |
| Machine_learning | Artificial_neural_network | 0.8773092031478882 |
| Machine_learning | Data_analysis | 0.8758628368377686 |
| Machine_learning | List_of_algorithms | 0.8737979531288147 |
| Machine_learning | K-means_clustering | 0.8715602159500122 |
| Albert_Einstein | Albert_Einstein | 1.0000001192092896 |
| Albert_Einstein | Physics | 0.8664291501045227 |
| Albert_Einstein | Werner_Heisenberg | 0.8625140190124512 |
| Albert_Einstein | Richard_Feynman | 0.8496938943862915 |
| Albert_Einstein | List_of_physicists | 0.8415523767471313 |
| Albert_Einstein | Physicist | 0.8384397625923157 |
| Albert_Einstein | Max_Planck | 0.8370327353477478 |
| Albert_Einstein | Niels_Bohr | 0.8340970873832703 |
| Albert_Einstein | Quantum_mechanics | 0.8331197500228882 |
| Albert_Einstein | Special_relativity | 0.8280861973762512 |
+-----+
```

## 8.1.8 Getting All Trained Vertex Vectors

You can retrieve the trained vertex vectors for the current DeepWalk model and store it in the database as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
opg4j> var vertexVectors = model.getTrainedVertexVectors().flattenAll()
opg4j> vertexVectors.write().db().name("vertex
vectors").tablename("vertexVectors").overwrite(true).store()
```

### Java

```
PgxFrame vertexVectors = model.getTrainedVertexVectors().flattenAll();
vertexVectors.write()
 .db()
 .name("vertex vectors")
 .tablename("vertexVectors")
 .overwrite(true)
 .store();
```

### Python

```
vertex_vectors = model.trained_vectors.flatten_all()
vertex_vectors.write().db().table_name("table_name").name("vertex_vectors").o
verwrite(True).store()
```

---

## 8.1.9 Storing a Trained DeepWalk Model

You can store models in database. The models get stored as a row inside a model store table.

The following code shows how to store a trained DeepWalk model in database in a specific model store table:

- 
- [JShell](#)
  - [Java](#)

- Python

## JShell

```
opg4j> model.export().db().
 modelstore("modelstoretablename"). // name of the model
store table
 modelname("model"). // model name
(primary key of model store table)
 description("a model description"). // description to
store alongside the model
 store()
```

## Java

```
model.export().db()
 .modelstore("modelstoretablename") // name of the model store
table
 .modelname("model") // model name (primary key of
model store table)
 .description("a model description") // description to store
alongside the model
 .store();
```

## Python

```
model.export().db(model_store="modeltablename", # name of
the model store table
 model_name="model", # model
name (primary key of model store table)
 model_description="a model description") #
description to store alongside the model
```



### Note:

All the preceding examples assume that you are storing the model in the current logged in database. If you must store the model in a different database then refer to the examples in [Storing a Trained Model in Another Database](#).

- [Storing a Trained Model in Another Database](#)

### 8.1.9.1 Storing a Trained Model in Another Database

You can store models in a different database other than the one used for login.

The following code shows how to store a trained model in a different database:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> model.export().db().
 username("user"). // DB user to use for
storing the model
 password("password"). // password of the DB user
 jdbcUrl("jdbcUrl"). // jdbc url to the DB
 modelstore("modelstoretablename"). // name of the model store
table
 modelName("model"). // model name (primary key
of model store table)
 description("a model description"). // description to store
alongside the model
 store()
```

## Java

```
model.export().db()
 .username("user") // DB user to use for storing the
model
 .password("password") // password of the DB user
 .jdbcUrl("jdbcUrl") // jdbc url to the DB
 .modelstore("modelstoretablename") // name of the model store table
 .modelName("model") // model name (primary key of model
store table)
 .description("a model description") // description to store alongside
the model
 .store();
```

## Python

```
model.export().db(username="user", # DB user to use
for storing the model
 password="password", # password of the
DB user
 jdbc_url="jdbc_url", # jdbc url to the
DB
 model_store="modelstoretablename", # name of the
model store table
 model_name="model", # model name
(primary key of model store table)
 model_description="a model description") # description to
store alongside the model
```

## 8.1.10 Loading a Pre-Trained DeepWalk Model

You can load models from a database.

You can load a pre-trained DeepWalk model from a model store table in database as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> var model = analyst.loadDeepWalkModel().db().
 modelstore("modeltablename"). // name of the model
store table
 modelname("model"). // model name (primary
key of model store table)
 load()
```

### Java

```
DeepWalkModel model = analyst.loadDeepWalkModel().db()
 .modelstore("modeltablename") // name of the model store table
 .modelname("model") // model name (primary key of model
store table)
 .load();
```

### Python

```
model = analyst.get_deepwalk_model_loader(). \
 db(model_store="modelstoretablename", # name of the
model store table
 model_name="model") # model name
(primary key of model store table)
```

 **Note:**

All the preceding examples assume that you are loading the model from the current logged in database. If you must load the model from a different database then refer to the examples in [Loading a Pre-Trained Model From Another Database](#).

- [Loading a Pre-Trained Model From Another Database](#)

### 8.1.10.1 Loading a Pre-Trained Model From Another Database

You can load models from a different database other than the one used for login.

You can load a pre-trained model from a model store table in database as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> var model = analyst.<modelLoader>.db().
 username("user"). // DB user to use for storing
the model
 password("password"). // password of the DB user
 jdbcUrl("jdbcUrl"). // jdbc url to the DB
 modelstore("modeltablename"). // name of the model store
table
 modelName("model"). // model name (primary key of
model store table)
 load()
```

where `<modelLoader>` applies as follows:

- `loadDeepWalkModel()`: Loads a Deepwalk model
- `loadSupervisedGraphWiseModel()`: Loads a Supervised GraphWise model
- `loadUnsupervisedGraphWiseModel()`: Loads a Unsupervised GraphWise model
- `loadPg2vecModel()`: Loads a Pg2vec model

#### Java

```
<modeltype> model = analyst.<modelLoader>.db()
 .username("user") // DB user to use for storing the model
 .password("password") // password of the DB user
 .jdbcUrl("jdbcUrl") // jdbc url to the DB
 .modelstore("modeltablename") // name of the model store table
 .modelName("model") // model name (primary key of model store
table)
 .load();
```

where `<modeltype>` can have the following values based on the model to be loaded:

- `DeepWalkModel`: represents a Deepwalk model
- `SupervisedGraphWiseModel`: represents a Supervised GraphWise model



- `UnsupervisedGraphWiseModel`: represents a Unsupervised GraphWise model
- `Pg2vecModel`: represents a Pg2vec model

where `<modelLoader>` applies as follows:

- `loadDeepWalkModel()`: Loads a Deepwalk model
- `loadSupervisedGraphWiseModel()`: Loads a Supervised GraphWise model
- `loadUnsupervisedGraphWiseModel()`: Loads a Unsupervised GraphWise model
- `loadPg2vecModel()`: Loads a Pg2vec model

## Python

```
model = analyst.<modelLoader>.db(username="user", # DB user to use
for storing the model
 password="password", # password of
the DB user
 jdbc_url="jdbc_url", # jdbc url to
the DB
 model_store="modelstoretablename", # name of the
model store table
 model_name="model") # model name
(primary key of model store table)
```

where `<modelLoader>` applies as follows:

- `get_deepwalk_model_loader()`: Loads a Deepwalk model
- `get_supervised_graphwise_model_loader()`: Loads a Supervised GraphWise model
- `get_unsupervised_graphwise_model_loader()`: Loads a Unsupervised GraphWise model
- `get_pg2vec_model_loader()`: Loads a Pg2vec model

---

### 8.1.11 Destroying a DeepWalk Model

You can destroy a DeepWalk model as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> model.destroy()
```

## Java

```
model.destroy();
```

## Python

```
model.destroy()
```

---

## 8.2 Using the Supervised GraphWise Algorithm

**Supervised GraphWise** is an inductive vertex representation learning algorithm which is able to leverage vertex feature information. It can be applied to a wide variety of tasks, including vertex classification and link prediction.

Supervised GraphWise is based on [GraphSAGE](#) by Hamilton et al.

### Model Structure

A Supervised GraphWise model consists of graph convolutional layers followed by several prediction layers.

The forward pass through a convolutional layer for a vertex proceeds as follows:

1. A set of neighbors of the vertex is sampled.
2. The previous layer representations of the neighbors are mean-aggregated, and the aggregated features are concatenated with the previous layer representation of the vertex.
3. This concatenated vector is multiplied with weights, and a bias vector is added.
4. The result is normalized to such that the layer output has unit norm.

The prediction layers are standard neural network layers.

The following describes the usage of the main functionalities of the implementation of **GraphSAGE** in PGX using the [Cora](#) graph as an example:

- [Loading a Graph](#)
- [Building a Minimal GraphWise Model](#)
- [Advanced Hyperparameter Customization](#)
- [Setting a Custom Loss Function and Batch Generator \(for Anomaly Detection\)](#)
- [Training a Supervised GraphWise Model](#)
- [Getting the Loss Value For a Supervised GraphWise Model](#)
- [Inferring the Vertex Labels for a Supervised GraphWise Model](#)
- [Evaluating the Supervised GraphWise Model Performance](#)
- [Inferring Embeddings for a Supervised GraphWise Model](#)
- [Storing a Trained Supervised GraphWise Model](#)
- [Loading a Pre-Trained Supervised GraphWise Model](#)

- [Destroying a Supervised GraphWise Model](#)
- [Explaining a Prediction of a Supervised GraphWise Model](#)

## 8.2.1 Loading a Graph

The following describes the steps for loading a graph:

1. Create a **Session** and an **Analyst**.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
cd /opt/oracle/graph/
./bin/opg4j
// starting the shell will create an implicit session and analyst
opg4j> import oracle.pgx.config.mllib.ActivationFunction
opg4j> import oracle.pgx.config.mllib.WeightInitScheme
opg4j> PgxSession session = Pgx.createSession("my-session")
opg4j> Analyst analyst = session.createAnalyst()
```

### Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.mllib.SupervisedGraphWiseModel;
import oracle.pgx.api.frames.*;
import oracle.pgx.config.mllib.ActivationFunction;
import oracle.pgx.config.mllib.GraphWiseConvLayerConfig;
import oracle.pgx.config.mllib.GraphWisePredictionLayerConfig;
import oracle.pgx.config.mllib.SupervisedGraphWiseModelConfig;
import oracle.pgx.config.mllib.WeightInitScheme;
PgxSession session = Pgx.createSession("my-session");
Analyst analyst = session.createAnalyst();
```

### Python

```
session = pypgx.get_session(session_name="my-session")
analyst = session.create_analyst()
```

- 
2. Load the **graph**.
- 

- [JShell](#)

- [Java](#)
- [Python](#)

## JShell

```
opg4j> var fullGraph = session.readGraphWithProperties("<path>/
<full_graph.json>")
opg4j> var trainGraph = session.readGraphWithProperties("<path>/
<train_graph.json>")
opg4j> var testVertices = fullGraph.getVertices()
 .stream()
 .filter(v -> !trainGraph.hasVertex(v.getId()))
 .collect(Collectors.toList())
```

## Java

```
PgxGraph fullGraph = session.readGraphWithProperties("<path>/
<full_graph.json>");
PgxGraph trainGraph = session.readGraphWithProperties("<path>/
<train_graph.json>");
List<PgxVertex> testVertices = fullGraph.getVertices()
 .stream()
 .filter(v->!trainGraph.hasVertex(v.getId()))
 .collect(Collectors.toList());
```

## Python

```
full_graph = session.read_graph_with_properties("<path>/
<full_graph.json>")
train_graph = session.read_graph_with_properties("<path>/
<train_graph.json>")
test_vertices = []
train_vertices = train_graph.get_vertices()
for v in full_graph.get_vertices():
 if not train_vertices.contains(v):
 test_vertices.append(v)
```

---

## 8.2.2 Building a Minimal GraphWise Model

You can build a GraphWise model using the minimal configuration and default hyper-parameters as described in the following code:

 **Note:**

Starting from Graph Server and Client Release 21.2, you can create a model with one of the following options:

- only vertex properties
- only edge properties
- both vertex and edge properties

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var model = analyst.supervisedGraphWiseModelBuilder().
 setVertexInputPropertyNames("features").
 setVertexTargetPropertyName("label").
 setEdgeInputPropertyNames("cost").
 build()
```

## Java

```
SupervisedGraphWiseModel model =
analyst.supervisedGraphWiseModelBuilder()
 .setVertexInputPropertyNames("features")
 .setVertexTargetPropertyName("labels")
 .setEdgeInputPropertyNames("cost")
 .build();
```

## Python

```
params = dict(vertex_target_property_name="label",
 edge_input_property_names=["cost"],
 vertex_input_property_names=["features"])

model = analyst.supervised_graphwise_builder(**params)
```

 **Note:**

Even though only one vertex and one edge property is specified in the preceding example, you can specify a list of vertex or edge properties.

## 8.2.3 Advanced Hyperparameter Customization

You can build a GraphWise model using rich hyperparameter customization.

This is done through the following two sub-config classes:

1. `GraphWiseConvLayerConfig`
2. `GraphWisePredictionLayerConfig`

The following code describes the implementation of the configuration using the above classes in GraphWise model:



### Note:

Starting from Graph Server and Client Release 21.2, you can create a model with one of the following options:

- only vertex properties
- only edge properties
- both vertex and edge properties

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> var weightProperty = analyst.pagerank(trainGraph).getName();
opg4j> var convLayerConfig = analyst.graphWiseConvLayerConfigBuilder().
 setNumSampledNeighbors(25).
 setActivationFunction(ActivationFunction.TANH).
 setWeightInitScheme(WeightInitScheme.XAVIER).
 setWeightedAggregationProperty(weightProperty).
 build()
opg4j> var predictionLayerConfig =
analyst.graphWisePredictionLayerConfigBuilder().
 setHiddenDimension(32).
 setActivationFunction(ActivationFunction.RELU).
 setWeightInitScheme(WeightInitScheme.HE).
 build()
opg4j> var model = analyst.supervisedGraphWiseModelBuilder().
 setVertexInputPropertyNames("features").
 setVertexTargetPropertyName("labels").
 setEdgeInputPropertyNames("cost").
 setConvLayerConfigs(convLayerConfig).
```

```
setPredictionLayerConfigs(predictionLayerConfig).
build()
```

## Java

```
String weightProperty = analyst.pagerank(trainGraph).getName()
GraphWiseConvLayerConfig convLayerConfig =
analyst.graphWiseConvLayerConfigBuilder()
 .setNumSampledNeighbors(25)
 .setActivationFunction(ActivationFunction.TANH)
 .setWeightInitScheme(WeightInitScheme.XAVIER)
 .setWeightedAggregationProperty(weightProperty)
 .build();

GraphWisePredictionLayerConfig predictionLayerConfig =
analyst.graphWisePredictionLayerConfigBuilder()
 .setHiddenDimension(32)
 .setActivationFunction(ActivationFunction.RELU)
 .setWeightInitScheme(WeightInitScheme.HE)
 .build();

SupervisedGraphWiseModel model =
analyst.supervisedGraphWiseModelBuilder()
 .setVertexInputPropertyNames("features")
 .setVertexTargetPropertyName("labels")
 .setEdgeInputPropertyNames("cost")
 .setConvLayerConfigs(convLayerConfig)
 .setPredictionLayerConfigs(predictionLayerConfig)
 .build();
```

## Python

```
weightProperty = analyst.pagerank(train_graph).name

conv_layer_config = dict(num_sampled_neighbors=25,
 activation_fn='TANH',
 weight_init_scheme='XAVIER',
 neighbor_weight_property_name=weightProperty)

conv_layer = analyst.graphwise_conv_layer_config(**conv_layer_config)

pred_layer_config = dict(hidden_dim=32,
 activation_fn='RELU',
 weight_init_scheme='HE')

pred_layer = analyst.graphwise_pred_layer_config(**pred_layer_config)

params = dict(vertex_target_property_name="labels",
 conv_layer_config=[conv_layer],
 pred_layer_config=[pred_layer],
 edge_input_property_names=["cost"],
 vertex_input_property_names=["features"],
 seed=17)
```

```
model = analyst.supervised_graphwise_builder(**params)
```

---

See [SupervisedGraphWiseModelBuilder](#), [GraphWiseConvLayerConfigBuilder](#) and [GraphWisePredictionLayerConfigBuilder](#) in Javadoc for a full description of all available hyperparameters and their default values.

## 8.2.4 Setting a Custom Loss Function and Batch Generator (for Anomaly Detection)

It is possible to select different loss functions for the supervised model by providing a `LossFunction` object, and different batch generators by providing a `BatchGenerator` object. This is useful for applications such as Anomaly Detection, which can be cast into the standard supervised framework but require different loss functions and batch generators.

`SupervisedGraphWise` model can use the [DevNetLoss](#) and the `StratifiedOversamplingBatchGenerator`. The `DevNetLoss` takes confidence margin and the value the anomaly takes in the target property as the two parameters.

The following example assumes that the `convLayerConfig` has already been defined:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
opg4j> import oracle.pgx.config.mllib.loss.LossFunctions
opg4j> import oracle.pgx.config.mllib.batchgenerator.BatchGenerators
opg4j> var predictionLayerConfig =
analyst.graphWisePredictionLayerConfigBuilder().
 setHiddenDimension(32).
 setActivationFunction(ActivationFunction.LINEAR).
 build()
opg4j> var model = analyst.supervisedGraphWiseModelBuilder().
 setVertexInputPropertyNames("vertex_features").
 setEdgeInputPropertyNames("edge_features").
 setVertexTargetPropertyName("labels").
 setConvLayerConfigs(convLayerConfig).
 setPredictionLayerConfigs(predictionLayerConfig).
 setLossFunction(LossFunctions.devNetLoss(5.0, true)).
 setBatchGenerator(BatchGenerators.STRATIFIED_OVERSAMPLING).
 build()
```



## Java

```
import oracle.pgx.config.mllib.loss.LossFunctions;
import oracle.pgx.config.mllib.batchgenerator.BatchGenerators;

GraphWisePredictionLayerConfig predictionLayerConfig =
 analyst.graphWisePredictionLayerConfigBuilder()
 .setHiddenDimension(32)
 .setActivationFunction(ActivationFunction.LINEAR)
 .build();

SupervisedGraphWiseModel model =
 analyst.supervisedGraphWiseModelBuilder()
 .setVertexInputPropertyNames("vertex_features")
 .setEdgeInputPropertyNames("edge_features")
 .setVertexTargetPropertyName("labels")
 .setConvLayerConfigs(convLayerConfig)
 .setPredictionLayerConfigs(predictionLayerConfig)
 .setLossFunction(LossFunctions.devNetLoss(5.0, true))
 .setBatchGenerator(BatchGenerators.STRATIFIED_OVERSAMPLING)
 .build();
```

## Python

```
from pypgx.api.mllib import DevNetLoss

pred_layer_config = dict(hidden_dim=32,
 activation_fn='LINEAR')

pred_layer = analyst.graphwise_pred_layer_config(**pred_layer_config)

params = dict(vertex_target_property_name="labels",
 conv_layer_config=[conv_layer],
 pred_layer_config=[pred_layer],
 vertex_input_property_names=["vertex_features"],
 edge_input_property_names=["edge_features"],
 loss_fn=DevNetLoss(5.0, True),
 batch_gen='Stratified_Oversampling',
 seed=17)

model = analyst.supervised_graphwise_builder(**params)
```

---

## 8.2.5 Training a Supervised GraphWise Model

You can train a Supervised GraphWise model on a graph as described in the following code:

- 
- [JShell](#)

- [Java](#)
- [Python](#)

## JShell

```
opg4j> model.fit(trainGraph)
```

## Java

```
model.fit(trainGraph);
```

## Python

```
model.fit(train_graph)
```

---

## 8.2.6 Getting the Loss Value For a Supervised GraphWise Model

You can fetch the training loss value as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var loss = model.getTrainingLoss()
```

## Java

```
double loss = model.getTrainingLoss();
```

## Python

```
loss = model.get_training_loss()
```

---

## 8.2.7 Inferring the Vertex Labels for a Supervised GraphWise Model

You can infer the labels for vertices on any graph (including vertices or graphs that were not seen during training) as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var labels = model.inferLabels(fullGraph, testVertices)
opg4j> labels.head().print()
```

## Java

```
PgxFrame labels = model.inferLabels(fullGraph, testVertices);
labels.head().print();
```

## Python

```
labels = model.infer_labels(full_graph, test_vertices)
labels.print()
```

The output will be similar to the following example output:

```
+-----+
| vertexId | label |
+-----+
| 2 | Neural Networks |
| 6 | Theory |
| 7 | Case Based |
| 22 | Rule Learning |
| 30 | Theory |
| 34 | Neural Networks |
| 47 | Case Based |
| 48 | Probabalistic Methods |
| 50 | Theory |
| 52 | Theory |
+-----+
```

Similarly, you can also get the model confidence for each class by inferring the prediction logits as described in the following code:

- [JShell](#)
- [Java](#)

- Python

## JShell

```
opg4j> var logits = model.inferLogits(fullGraph, testVertices)
opg4j> logits.head().print()
```

## Java

```
PgxFrame logits = model.inferLogits(fullGraph, testVertices);
logits.head().print();
```

## Python

```
logits = model.infer_logits(full_graph, test_vertices)
logits.print()
```

---

## 8.2.8 Evaluating the Supervised GraphWise Model Performance

You can evaluate various classification metrics for the model using the `evaluateLabels` method as described in the following code:

- 
- JShell
  - Java
  - Python

## JShell

```
opg4j> model.evaluateLabels(fullGraph, testVertices).print()
```

## Java

```
model.evaluateLabels(fullGraph, testVertices).print();
```

## Python

```
model.evaluate_labels(full_graph, test_vertices).print()
```

---

The output will be similar to the following example output:

```
+-----+
| Accuracy | Precision | Recall | F1-Score |
+-----+
| 0.8488 | 0.8523 | 0.831 | 0.8367 |
+-----+
```

## 8.2.9 Inferring Embeddings for a Supervised GraphWise Model

You can use a trained model to infer embeddings for unseen nodes and store in the database as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> var vertexVectors = model.inferEmbeddings(fullGraph,
testVertices).flattenAll()
opg4j> vertexVectors.write().
 db().
 name("vertex vectors").
 tablename("vertexVectors").
 overwrite(true).
 store()
```

### Java

```
PgxFrame vertexVectors =
model.inferEmbeddings(fullGraph,testVertices).flattenAll();
vertexVectors.write()
 .db()
 .name("vertex vectors")
 .tablename("vertexVectors")
 .overwrite(true)
 .store();
```

### Python

```
vertex_vectors = model.infer_embeddings(full_graph,
test_vertices).flatten_all()
vertex_vectors.write().db().table_name("table_name").name("vertex_vecto
rs").overwrite(True).store()
```

The schema for the `vertexVectors` will be as follows without flattening (`flattenAll` splits the vector column into separate double-valued columns):

```
+-----+
| vertexId | embedding |
+-----+
```



**Note:**

All the preceding examples assume that you are inferring the embeddings for a model in the current logged in database. If you must infer embeddings for the model in a different database then refer to the examples in [Inferring Embeddings for a Model in Another Database](#).

- [Inferring Embeddings for a Model in Another Database](#)

### 8.2.9.1 Inferring Embeddings for a Model in Another Database

You can infer embeddings on a trained model and store in a different database other than the one used for login.

The following code shows how to infer embeddings and store in a different database:

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg-jshell> var vertexVectors = model.inferEmbeddings(fullGraph,
testVertices).flattenAll()
opg-jshell> vertexVectors.write().
 db().
 username("user"). // DB user to use for storing the
model
 password("password"). // password of the DB user
 jdbcUrl("jdbcUrl"). // jdbc url to the DB
 name("vertex vectors").
 tablename("vertexVectors"). // indicates the name of the table
in which the data should be stored
 overwrite(true).
 store()
```

## Java

```
PgxFrame vertexVectors =
model.inferEmbeddings(fullGraph,testVertices).flattenAll();
vertexVectors.write()
 .db()
 .username("user") // DB user to use for storing
the model
 .password("password") // password of the DB user
 .jdbcUrl("jdbcUrl") // jdbc url to the DB
 .name("vertex vectors")
 .tablename("vertexVectors") // indicates the name of the
table in which the data should be stored
 .overwrite(true)
 .store();
```

## Python

```
vertex_vectors =
model.infer_embeddings(fullGraph,test_vertices).flattenAll()
vertex_vectors.write().db().username("user") \
 .password("password") \
 .jdbc_url("jdbcUrl") \
 .table_name("table_name") \
 .name("vertex vectors") \
 .overwrite(True) \
 .store()
```

---

### 8.2.10 Storing a Trained Supervised GraphWise Model

You can store models in database. The models get stored as a row inside a model store table.

The following code shows how to store a trained Supervised GraphWise model in database in a specific model store table:

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg4j> model.export().db().
 modelstore("modelstoretablename"). // name of the model
store table
 modelname("model"). // model name
(primary key of model store table)
```

```
description("a model description"). // description to store
alongside the model
store()
```

## Java

```
model.export().db()
 .modelstore("modelstoretablename") // name of the model store table
 .modelname("model") // model name (primary key of model
store table)
 .description("a model description") // description to store alongside
the model
 .store();
```

## Python

```
model.export().db(model_store="modeltablename", # name of the
model store table
 model_name="model", # model name
(primary key of model store table)
 model_description="a model description") # description to
store alongside the model
```

### Note:

All the preceding examples assume that you are storing the model in the current logged in database. If you must store the model in a different database then refer to the examples in [Storing a Trained Model in Another Database](#).

## 8.2.11 Loading a Pre-Trained Supervised GraphWise Model

You can load models from a database.

You can load a pre-trained Supervised GraphWise model from a model store table in database as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> var model = analyst.loadSupervisedGraphWiseModel().db().
modelstore("modeltablename"). // name of the model store
```



```
table
 modelname("model"). // model name (primary
key of model store table)
 load();
```

## Java

```
SupervisedGraphWiseModel model =
analyst.loadSupervisedGraphWiseModel().db()
 .modelstore("modeltablename") // name of the model store table
 .modelname("model") // model name (primary key of model
store table)
 .load();
```

## Python

```
model = analyst.get_supervised_graphwise_model_loader(). \
 db(model_store="modelstoretablename", # name of the
model store table
 model_name="model") # model name
(primary key of model store table)
```



### Note:

All the preceding examples assume that you are loading the model from the current logged in database. If you must load the model from a different database then refer to the examples in [Loading a Pre-Trained Model From Another Database](#).

## 8.2.12 Destroying a Supervised GraphWise Model

You can destroy a GraphWise model as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> model.destroy()
```

## Java

```
model.destroy();
```

## Python

```
model.destroy()
```

---

### 8.2.13 Explaining a Prediction of a Supervised GraphWise Model

In order to understand which features and vertices are important for a prediction of the Supervised GraphWise model, you can generate a `SupervisedGnnExplanation` using a technique similar to the [GNNE explainer](#) by Ying et al.

The explanation holds information related to:

- **graph structure:** an importance score for each vertex
- **features:** an importance score for each graph property

 **Note:**

The vertex being explained is always assigned importance 1. Further, the feature importances are scaled such that the most important feature has importance 1.

Additionally, a `SupervisedGnnExplanation` contains the inferred embedding, logits, and label. The `inferAndGetExplanation` method can be used on all fitted `SupervisedGraphWiseModel` models that do not rely on edge features. In order to achieve best results, the features should be centered around 0.

For example, assume a simple graph that contains a feature that correlates with the label and another feature that does not. It is therefore expected that the importance of the features to differ significantly (with the feature correlating with the label being more important), while structural importance does not play a big role. In this case, you can generate an explanation as shown:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> var simpleGraph = session.createGraphBuilder().
 addVertex(0).setProperty("label_feature",
0.5).setProperty("const_feature", 0.5).
 setProperty("label", true).
```

```

 addVertex(1).setProperty("label_feature",
-0.5).setProperty("const_feature", 0.5).
 setProperty("label", false).
 addEdge(0, 1).build()

// build and train a Supervised GraphWise model

// explain prediction of vertex 0
opg4j> var explanation = model.inferAndGetExplanation(simpleGraph,
simpleGraph.getVertex(0))
opg4j> var constProperty =
simpleGraph.getVertexProperty("const_feature")
opg4j> var labelProperty =
simpleGraph.getVertexProperty("label_feature")

// retrieve feature importances
opg4j> var featureImportances =
explanation.getVertexFeatureImportance()
opg4j> var importanceConstProp =
featureImportances.get(constProperty) // small as unimportant
opg4j> var importanceLabelProp =
featureImportances.get(labelProperty) // large (1) as important

// retrieve computation graph with importances
opg4j> var importanceGraph = explanation.getImportanceGraph()

// retrieve importance of vertices
opg4j> var importanceProperty =
explanation.getVertexImportanceProperty()
opg4j> var importanceVertex0 = importanceProperty.get(0) // has
importance 1
opg4j> var importanceVertex1 = importanceProperty.get(1) // available
if vertex 1 part of computation

```

## Java

```

PgxGraph simpleGraph = session.createGraphBuilder()
 .addVertex(0).setProperty("label_feature",
0.5).setProperty("const_feature", 0.5)
 .setProperty("label", true)
 .addVertex(1).setProperty("label_feature",
-0.5).setProperty("const_feature", 0.5)
 .setProperty("label", false)
 .addEdge(0, 1).build();

// build and train a Supervised GraphWise model

// explain prediction of vertex 0
SupervisedGnnExplanation<Integer> explanation =
model.inferAndGetExplanation(simpleGraph,
 simpleGraph.getVertex(0));

VertexProperty<Integer, Float> constProperty =

```

```
simpleGraph.getVertexProperty("const_feature");
VertexProperty<Integer, Float> labelProperty =
simpleGraph.getVertexProperty("label_feature");

// retrieve feature importances
Map<VertexProperty<Integer, ?>, Float> featureImportances =
explanation.getVertexFeatureImportance();
float importanceConstProp = featureImportances.get(constProperty); // small
as unimportant
float importanceLabelProp = featureImportances.get(labelProperty); // large
(1) as important

// retrieve computation graph with importances
PgxGraph importanceGraph = explanation.getImportanceGraph();

// retrieve importance of vertices
VertexProperty<Integer, Float> importanceProperty =
explanation.getVertexImportanceProperty();
float importanceVertex0 = importanceProperty.get(0); // has importance 1
float importanceVertex1 = importanceProperty.get(1); // available if vertex
1 part of computation
```

## Python

```
simple_graph = session.create_graph_builder()
 .add_vertex(0).set_property("label_feature",
0.5).set_property("const_feature", 0.5)
 .set_property("label", true)
 .add_vertex(1).set_property("label_feature",
-0.5).set_property("const_feature", 0.5)
 .set_property("label", false)
 .add_edge(0, 1).build()

build and train a Supervised GraphWise model

explain prediction of vertex 0
explanation = model.infer_and_get_explanation(simple_graph,
simple_graph.get_vertex(0))

const_property = simple_graph.get_vertex_property("const_feature")
label_property = simple_graph.get_vertex_property("label_feature")

retrieve feature importances
feature_importances = explanation.get_vertex_feature_importance()
importance_const_prop = feature_importances[const_property]
importance_label_prop = feature_importances[label_property]

retrieve computation graph with importances
importance_graph = explanation.get_importance_graph()

retrieve importance of vertices
importance_property = explanation.get_vertex_importance_property()
importance_vertex_0 = importance_property[0]
importance_vertex_1 = importance_property[1]
```

 **See Also:**

- [Building a Minimal GraphWise Model](#)
- [Training a Supervised GraphWise Model](#)

## 8.3 Using the Unsupervised GraphWise Algorithm

**Unsupervised GraphWise** is an unsupervised inductive vertex representation learning algorithm which is able to leverage vertex information. The learned embeddings can be used in various downstream tasks including vertex classification, vertex clustering and similar vertex search.

Unsupervised GraphWise is based on [Deep Graph Infomax \(DGI\)](#) by Velickovic et al.

### Model Structure

A Unsupervised GraphWise model consists of graph convolutional layers followed by a DGI Layer.

The forward pass through a convolutional layer for a vertex proceeds as follows:

1. A set of neighbors of the vertex is sampled.
2. The previous layer representations of the neighbors are mean-aggregated, and the aggregated features are concatenated with the previous layer representation of the vertex.
3. This concatenated vector is multiplied with weights, and a bias vector is added.
4. The result is normalized to such that the layer output has unit norm.

The DGI Layer consists of three parts enabling unsupervised learning using embeddings produced by the convolution layers.

1. **Corruption function:** Shuffles the node features while preserving the graph structure to produce negative embedding samples using the convolution layers.
2. **Readout function:** Sigmoid activated mean of embeddings, used as summary of a graph.
3. **Discriminator:** Measures the similarity of positive (unshuffled) embeddings with the summary as well as the similarity of negative samples with the summary from which the loss function is computed.

Since none of these contains mutable hyperparameters, the default DGI layer is always used and cannot be adjusted.

The following describes the usage of the main functionalities of the implementation of **DGI** in PGX using the [Cora](#) graph as an example:

- [Loading a Graph](#)
- [Building a Minimal Unsupervised GraphWise Model](#)

- [Advanced Hyperparameter Customization](#)
- [Training a Unsupervised GraphWise Model](#)
- [Getting the Loss Value for a Unsupervised GraphWise Model](#)
- [Inferring Embeddings for a Unsupervised GraphWise Model](#)
- [Storing a Unsupervised GraphWise Model](#)
- [Loading a Pre-Trained Unsupervised GraphWise Model](#)
- [Destroying a Unsupervised GraphWise Model](#)
- [Explaining a Prediction for an Unsupervised GraphWise Model](#)

## 8.3.1 Loading a Graph

The following describes the steps for loading a graph:

1. Create a **Session** and an **Analyst**.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
cd /opt/oracle/graph/
./bin/opg4j
// starting the shell will create an implicit session and analyst
opg4j> import oracle.pgx.config.mllib.ActivationFunction
opg4j> import oracle.pgx.config.mllib.WeightInitScheme
opg4j> PgxSession session = Pgx.createSession("my-session")
opg4j> Analyst analyst = session.createAnalyst()
```

### Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.mllib.UnsupervisedGraphWiseModel;
import oracle.pgx.api.frames.*;
import oracle.pgx.config.mllib.ActivationFunction;
import oracle.pgx.config.mllib.GraphWiseConvLayerConfig;
import oracle.pgx.config.mllib.UnsupervisedGraphWiseModelConfig;
import oracle.pgx.config.mllib.WeightInitScheme;

PgxSession session = Pgx.createSession("my-session");
Analyst analyst = session.createAnalyst();
```

## Python

```
session = pypgx.get_session()
analyst = session.analyst
```

---

### 2. Load the **graph**.

---

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var graph = session.readGraphWithProperties("<path/to/
graph_config.json>")
```

## Java

```
PgxGraph graph = session.readGraphWithProperties("<path/to/
graph_config.json>");
```

## Python

```
graph = session.read_graph_with_properties("<path/to/
graph_config.json>")
```

---

You do not need to use a test graph or test vertices, since the model is trained to be unsupervised.

## 8.3.2 Building a Minimal Unsupervised GraphWise Model

You can build a Unsupervised GraphWise model with only vertex properties, or only edge properties or both using the minimal configuration and default hyper-parameters.

---

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var model = analyst.unsupervisedGraphWiseModelBuilder().
 setVertexInputPropertyNames("features").
 build()
```

## Java

```
UnsupervisedGraphWiseModel model =
analyst.unsupervisedGraphWiseModelBuilder()
 .setVertexInputPropertyNames("features")
 .build();
```

## Python

```
model =
analyst.unsupervised_graphwise_builder(vertex_input_property_names=["features
"])
```

---

### 8.3.3 Advanced Hyperparameter Customization

You can build a Unsupervised GraphWise model with only vertex properties or only edge properties or both using rich hyperparameter customization.

This is implemented using the sub-config class, `GraphWiseConvLayerConfig`.

The following code describes the implementation of the configuration in a Unsupervised GraphWise model:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var weightProperty = analyst.pagerank(trainGraph).getName()
opg4j> var convLayerConfig = analyst.graphWiseConvLayerConfigBuilder().
 setNumSampledNeighbors(25).
 setActivationFunction(ActivationFunction.TANH).
 setWeightInitScheme(WeightInitScheme.XAVIER).
 setWeightedAggregationProperty(weightProperty).
 build()

opg4j> var model = analyst.unsupervisedGraphWiseModelBuilder().
 setVertexInputPropertyNames("features").
```



```
setConvLayerConfigs(convLayerConfig).
build()
```

## Java

```
String weightProperty = analyst.pagerank(trainGraph).getName();
GraphWiseConvLayerConfig convLayerConfig =
analyst.graphWiseConvLayerConfigBuilder()
 .setNumSampledNeighbors(25)
 .setActivationFunction(ActivationFunction.TANH)
 .setWeightInitScheme(WeightInitScheme.XAVIER)
 .setWeightedAggregationProperty(weightProperty)
 .build();

UnsupervisedGraphWiseModel model =
analyst.unsupervisedGraphWiseModelBuilder()
 .setVertexInputPropertyNames("features")
 .setConvLayerConfigs(convLayerConfig)
 .build();
```

## Python

```
weightProperty = analyst.pagerank(train_graph).name

conv_layer_config = dict(num_sampled_neighbors=25,
 activation_fn='TANH',
 weight_init_scheme='XAVIER',
 neighbor_weight_property_name=weightProperty)
conv_layer = analyst.graphwise_conv_layer_config(**conv_layer_config)
params = dict(conv_layer_config=[conv_layer],
 vertex_input_property_names=["features"])

model = analyst.unsupervised_graphwise_builder(**params)
```

---

See [UnsupervisedGraphWiseModelBuilder](#) and [GraphWiseConvLayerConfigBuilder](#) in Javadoc for full description of all available hyperparameters and their default values.

### 8.3.4 Training a Unsupervised GraphWise Model

You can train a Unsupervised GraphWise model on a graph as shown:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> model.fit(trainGraph)
```

## Java

```
model.fit(trainGraph);
```

## Python

```
model.fit(train_graph)
```

---

### 8.3.5 Getting the Loss Value for a Unsupervised GraphWise Model

You can fetch the training loss value for a Unsupervised GraphWise Model as shown in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var loss = model.getTrainingLoss()
```

## Java

```
double loss = model.getTrainingLoss();
```

## Python

```
loss = model.get_training_loss()
```

---

### 8.3.6 Inferring Embeddings for a Unsupervised GraphWise Model

You can use a trained model to infer embeddings for unseen nodes and store them in the database as described in the following code:

---

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var vertexVectors = model.inferEmbeddings(fullGraph,
fullGraph.getVertices()).flattenAll()
opg4j> vertexVectors.write().
 db().
 name("vertex vectors").
 tablename("vertexVectors").
 overwrite(true).
 store()
```

## Java

```
PgxFrame vertexVectors =
model.inferEmbeddings(fullGraph,fullGraph.getVertices()).flattenAll();
vertexVectors.write()
 .db()
 .name("vertex vectors")
 .tablename("vertexVectors")
 .overwrite(true)
 .store();
```

## Python

```
vertex_vectors =
model.infer_embeddings(full_Graph,full_Graph.get_vertices()).flattenAl
l()
vertex_vectors.write().db().table_name("table_name").name("vertex_vecto
rs").overwrite(True).store()
```

---

The schema for the `vertexVectors` will be as follows without flattening (`flattenAll` splits the vector column into separate double-valued columns):

```
+-----+
| vertexId | embedding |
+-----+
```

 **Note:**

All the preceding examples assume that you are inferring the embeddings for a model in the current logged in database. If you must infer embeddings for the model in a different database then refer to the examples in [Inferring Embeddings for a Model in Another Database](#).

## 8.3.7 Storing a Unsupervised GraphWise Model

You can store models in database. The models get stored as a row inside a model store table.

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> model.export().db().
 modelstore("modelstoretablename"). // name of the model store
table
 modelname("model"). // model name (primary key
of model store table)
 description("a model description"). // description to store
alongside the model
 store()
```

### Java

```
model.export().db()
 .modelstore("modelstoretablename") // name of the model store table
 .modelname("model") // model name (primary key of model
store table)
 .description("a model description") // description to store alongside
the model
 .store();
```

### Python

```
model.export().db(model_store="modeltablename", # name of the
model store table
 model_name="model", # model name
(primary key of model store table)
 model_description="a model description") # description to
store alongside the model
```

**Note:**

All the preceding examples assume that you are storing the model in the current logged in database. If you must store the model in a different database then refer to the examples in [Storing a Trained Model in Another Database](#).

## 8.3.8 Loading a Pre-Trained Unsupervised GraphWise Model

You can load models from a database.

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> var model = analyst.loadUnsupervisedGraphWiseModel().db().
 modelstore("modeltablename"). // name of the model
store table
 modelname("model"). // model name (primary
key of model store table)
 load()
```

### Java

```
UnsupervisedGraphWiseModel model =
analyst.loadUnsupervisedGraphWiseModel().db()
 .modelstore("modeltablename") // name of the model store table
 .modelname("model") // model name (primary key of model
store table)
 .load();
```

### Python

```
model = analyst.get_unsupervised_graphwise_model_loader(). \
 db(model_store="modelstoretablename", # name of the
model store table
 model_name="model") # model name
(primary key of model store table)
```

 **Note:**

All the preceding examples assume that you are loading the model from the current logged in database. If you must load the model from a different database then refer to the examples in [Loading a Pre-Trained Model From Another Database](#).

### 8.3.9 Destroying a Unsupervised GraphWise Model

You can destroy a Unsupervised GraphWise model as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg-jshell> model.destroy()
```

#### Java

```
model.destroy();
```

#### Python

```
model.destroy()
```

---

### 8.3.10 Explaining a Prediction for an Unsupervised GraphWise Model

In order to understand which features and vertices are important for a prediction of the Unsupervised GraphWise model, you can generate an `UnsupervisedGnnExplanation` using a technique similar to the [GNNEExplainer](#) by Ying et al.

The explanation holds information related to:

- **graph structure**: an importance score for each vertex
- **features**: an importance score for each graph property

 **Note:**

The vertex being explained is always assigned importance 1. Further, the feature importances are scaled such that the most important feature has importance 1.

Additionally, an `UnsupervisedGnnExplanation` contains the inferred embedding. The `inferAndGetExplanation` method can be used on all fitted `UnsupervisedGraphWiseModel` models. In order to achieve best results, the features should be centered around 0.

For example, assume a simple graph, `componentGraph` which contains `k` densely connect *components*, that is, there are many edges between vertices of the same component and few edges between any two components. By training an Unsupervised GraphWise model on this graph, you can expect a model that produces similar embeddings for vertices in a densely connected component.

The following example shows how to generate an explanation on an inference `componentGraph`. It is expected that vertices from the same component to have a higher importance than vertices from a different component. Note that the feature importances are not relevant in this example.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var componentGraph =
session.readGraphWithProperties("<path_to_component_graph.json>")
// explain prediction of vertex 0
opg4j> var feat1Property = componentGraph.getVertexProperty("feat1")
opg4j> var feat2Property = componentGraph.getVertexProperty("feat2")

// build and train an Unsupervised GraphWise model

// explain prediction of vertex 0
// setting the numClusters argument to the expected number of clusters
may improve
// explanation results
opg4j> var explanation = model.inferAndGetExplanation(componentGraph,
componentGraph.getVertex(0), 10)

// retrieve computation graph with importance
opg4j> var importanceGraph = explanation.getImportanceGraph()

// retrieve importance of vertices
// vertex 1 is in the same densely connected component as vertex 0
// vertex 2 is in a different component
opg4j> var importanceProperty =
explanation.getVertexImportanceProperty()
opg4j> var importanceVertex0 = importanceProperty.get(0) // has
importance 1
opg4j> var importanceVertex1 = importanceProperty.get(1) // high
importance
opg4j> var importanceVertex2 = importanceProperty.get(2) // low
importance
```

```
opg4j> var featureImportances = explanation.getVertexFeatureImportance()
opg4j> var importanceConstProp = featureImportances.get(constProperty) //
small as unimportant
opg4j> var importanceLabelProp = featureImportances.get(labelProperty) //
large (1) as important

// optionally retrieve feature importance
opg4j> var featureImportances = explanation.getVertexFeatureImportance()
opg4j> var importanceFeat1Prop = featureImportances.get(featlProperty)
opg4j> var importanceFeat2Prop = featureImportances.get(featl2Property)
```

## Java

```
PgxGraph componentGraph =
session.readGraphWithProperties("<path_to_component_graph.json>") // load
component graph
VertexProperty<Integer, Float> feat1Property =
componentGraph.getVertexProperty("feat1");
VertexProperty<Integer, Float> feat2Property =
componentGraph.getVertexProperty("feat2");

// build and train an Unsupervised GraphWise model

// explain prediction of vertex 0
// setting the numClusters argument to the expected number of clusters may
improve
// explanation results
UnsupervisedGnnExplanation<Integer> explanation =
model.inferAndGetExplanation(componentGraph, componentGraph.getVertex(0),
10);

// retrieve computation graph with importances
PgxGraph importanceGraph = explanation.getImportanceGraph();

// retrieve importance of vertices
// vertex 1 is in the same densely connected component as vertex 0
// vertex 2 is in a different component
VertexProperty<Integer, Float> importanceProperty =
explanation.getVertexImportanceProperty();
float importanceVertex0 = importanceProperty.get(0); // has importance 1
float importanceVertex1 = importanceProperty.get(1); // high importance
float importanceVertex2 = importanceProperty.get(2); // low importance

// retrieve feature importance (not relevant for this example)
Map<VertexProperty<Integer, ?>, Float> featureImportances =
explanation.getVertexFeatureImportance();
float importanceFeat1Prop = featureImportances.get(featlProperty);
float importanceFeat2Prop = featureImportances.get(featl2Property);
```



## Python

```
load 'component_graph' with vertex features 'feat1' and 'feat2'
feat1_property = component_graph.get_vertex_property("feat1")
feat2_property = component_graph.get_vertex_property("feat2")

build and train an Unsupervised GraphWise model

explain prediction of vertex 0
setting the num_clusters argument to the expected number of clusters
may improve
explanation results
explanation = model.infer_and_get_explanation(
 graph=component_graph,
 vertex=component_graph.get_vertex(0),
 num_clusters=10,
)

retrieve computation graph with importances
importance_graph = explanation.get_importance_graph()

retrieve importance of vertices
vertex 1 is in the same densely connected component as vertex 0
vertex 2 is in a different component
importance_property = explanation.get_vertex_importance_property()
importance_vertex_0 = importance_property[0] # has importance 1
importance_vertex_1 = importance_property[1] # high importance
importance_vertex_2 = importance_property[2] # low importance

retrieve feature importance (not relevant for this example)
feature_importances = explanation.get_vertex_feature_importance()
importance_feat1_prop = feature_importances[feat1_property]
importance_feat2_prop = feature_importances[feat2_property]
```



### See Also:

- [Building a Minimal Unsupervised GraphWise Model](#)
- [Training a Unsupervised GraphWise Model](#)

## 8.4 Using the Pg2vec Algorithm

**Pg2vec** learns representations of graphlets (partitions inside a graph) by employing edges as the principal learning units and thereby packing more information in each learning unit (as compared to employing vertices as learning units) for the representation learning task.

It consists of three main steps:

1. Random walks for each vertex (with pre-defined length per walk and pre-defined number of walks per vertex) are generated.
2. Each edge in this random walk is mapped as a `property.edge-word` in the created document (with the document label as the graph-id) where the `property.edge-word` is defined as the concatenation of the properties of the source and destination vertices.
3. The generated documents (with their attached document labels) are fed to a `doc2vec` algorithm which generates the vector representation for each document, which is a graph in this case.

Pg2vec creates graphlet embeddings for a specific set of graphlets and cannot be updated to incorporate modifications on these graphlets. Instead, a new Pg2vec model should be trained on these modified graphlets.

The following represents the memory consumption of Pg2vec model.

$$O(2(n+m) * d)$$

where:

- `n`: is the number of vertices in the graph
- `m`: is the number of graphlets in the graph
- `d`: is the embedding length

The following describes the usage of the main functionalities of the implementation of Pg2vec in PGX using NCI109 dataset as an example with 4127 graphs in it:

- [Loading a Graph](#)
- [Building a Minimal Pg2vec Model](#)
- [Building a Customized Pg2vec Model](#)
- [Training a Pg2vec Model](#)
- [Getting the Loss Value For a Pg2vec Model](#)
- [Computing Similar Graphlets for a Given Graphlet](#)
- [Computing Similar for a Graphlet Batch](#)
- [Inferring a Graphlet Vector](#)
- [Inferring Vectors for a Graphlet Batch](#)
- [Storing a Trained Pg2vec Model](#)
- [Loading a Pre-Trained Pg2vec Model](#)
- [Destroying a Pg2vec Model](#)

## 8.4.1 Loading a Graph

The following describes the steps for loading a graph:

1. Create a **Session** and an **Analyst**.

- 
- [JShell](#)
  - [Java](#)

- [Python](#)

## JShell

```
cd /opt/oracle/graph/
./bin/opg4j
// starting the shell will create an implicit session and analyst
```

## Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.mllib.Pg2vecModel;
import oracle.pgx.api.frames.*;
...
PgxSession session = Pgx.createSession("my-session");
Analyst analyst = session.createAnalyst();
```

## Python

```
session = pypgx.get_session(session_name="my-session")
analyst = session.create_analyst()
```

---

## 2. Load the **graph**.

---

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var graph = session.readGraphWithProperties("<path>/
<graph.json>")
```

## Java

```
PgxGraph graph = session.readGraphWithProperties("<path>/
<graph.json>");
```

## Python

```
graph = session.read_graph_with_properties("<path>/<graph.json>")
```

---

## 8.4.2 Building a Minimal Pg2vec Model

You can build a Pg2vec model using the minimal configuration and default hyper-parameters as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> var model = analyst.pg2vecModelBuilder().
 setGraphLetIdPropertyName("graph_id").
 setVertexPropertyNames(Arrays.asList("category")).
 setWindowSize(4).
 setWalksPerVertex(5).
 setWalkLength(8).
 build()
```

### Java

```
Pg2vecModel model = analyst.pg2vecModelBuilder()
 .setGraphLetIdPropertyName("graph_id")
 .setVertexPropertyNames(Arrays.asList("category"))
 .setWindowSize(4)
 .setWalksPerVertex(5)
 .setWalkLength(8)
 .build();
```

### Python

```
model = analyst.pg2vec_model_builder(
 graph_let_id_property_name="graph_id",
 vertex_property_names(["category"]),
 window_size=4,
 walks_per_vertex=5,
 walk_length=8)
```

---

You can specify the property name to determine each graphlet using the `Pg2vecModelBuilder#setGraphLetIdPropertyName` operation and also employ the vertex properties in Pg2vec which are specified using the `Pg2vecModelBuilder#setVertexPropertyNames` operation.

You can also use the weakly connected component (WCC) functionality in PGX to determine the graphlets in a given graph.

## 8.4.3 Building a Customized Pg2vec Model

You can build a Pg2vec model using customized hyper-parameters as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
opg4j> var model = analyst.pg2vecModelBuilder().
 setGraphLetIdPropertyName("graph_id").
 setVertexPropertyNames(Arrays.asList("category")).
 setMinWordFrequency(1).
 setBatchSize(128).
 setNumEpochs(5).
 setLayerSize(200).
 setLearningRate(0.04).
 setMinLearningRate(0.0001).
 setWindowSize(4).
 setWalksPerVertex(5).
 setWalkLength(8).
 setUseGraphletSize(true).
 setValidationFraction(0.05).
 setGraphletSizePropertyName("<propertyName>").
 build()
```

### Java

```
Pg2vecModel model= analyst.pg2vecModelBuilder()
 .setGraphLetIdPropertyName("graph_id")
 .setVertexPropertyNames(Arrays.asList("category"))
 .setMinWordFrequency(1)
 .setBatchSize(128)
 .setNumEpochs(5)
 .setLayerSize(200)
 .setLearningRate(0.04)
 .setMinLearningRate(0.0001)
 .setWindowSize(4)
 .setWalksPerVertex(5)
 .setWalkLength(8)
 .setUseGraphletSize(true)
 .setValidationFraction(0.05)
 .setGraphletSizePropertyName("<propertyName>")
 .build();
```

## Python

```
model = analyst.pg2vec_model_builder(
 graph_id_property_name = "graph_id",
 vertex_property_names = ["category"],
 min_word_frequency = 1,
 batch_size = 128,
 num_epochs = 5,
 layer_size = 200,
 learning_rate = 0.04,
 min_learning_rate = 0.0001,
 window_size = 4,
 walks_per_vertex = 5,
 walk_length = 8,
 use_graphlet_size = true,
 graphlet_size_property_name = "<property_name>",
 validation_fraction = 0.05)
```

---

See [Pg2vecModelBuilder](#) in Javadoc for more explanation for each builder operation along with the default values.

### 8.4.4 Training a Pg2vec Model

You can train a Pg2vec model with the specified default or customized settings as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> model.fit(graph)
```

#### Java

```
model.fit(graph);
```

#### Python

```
model.fit(graph)
```

---

## 8.4.5 Getting the Loss Value For a Pg2vec Model

You can fetch the training loss value on a specified fraction of training data (set in builder using `setValidationFraction`) as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
opg4j> var loss = model.getLoss()
```

### Java

```
double loss = model.getLoss();
```

### Python

```
loss = model.loss
```

---

## 8.4.6 Computing Similar Graphlets for a Given Graphlet

You can fetch the  $k$  most similar graphlets for a given graphlet as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
opg4j> var similars = model.computeSimilars(52, 10)
```

### Java

```
PgxFrame similars = model.computeSimilars(52, 10);
```

---

## Python

```
similar = model.compute_similar(52, 10)
```

---


Searching for similar vertices for graphlet with ID = 52 using the trained model and printing it with `similar.print()`, will result in the following output:

```
+-----+
| dstGraphlet | similarity |
+-----+
| 52 | 1.0 |
| 10 | 0.8748674392700195 |
| 23 | 0.8551455140113831 |
| 26 | 0.8493421673774719 |
| 47 | 0.8411962985992432 |
| 25 | 0.8281504511833191 |
| 43 | 0.8202780485153198 |
| 24 | 0.8179885745048523 |
| 8 | 0.796689510345459 |
| 9 | 0.7947834134101868 |
+-----+
```

The following depicts the visualization of two similar graphlets (top: ID = 52 and bottom: ID = 10):



Figure 8-1 Pg2vec - Visualization of Two Similar Graphlets



## 8.4.7 Computing Similar for a Graphlet Batch

You can fetch the  $k$  most similar graphlets for a batch of input graphlets as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> var graphlets = new ArrayList()
opg4j> graphlets.add(52)
opg4j> graphlets.add(41)
opg4j> var batchedSimilar = model.computeSimilar(graphlets, 10)
```

### Java

```
List graphlets = Arrays.asList(52,41);
PgxFramework batchedSimilar = model.computeSimilar(graphlets,10);
```

## Python

```
batched_similars = model.compute_similars([52,41],10)
```

Searching for similar vertices for graphlet with ID = 52 and ID = 41 using the trained model and printing it with `batched_similars.print()`, will result in the following output:

```
+-----+
| srcGraphlet | dstGraphlet | similarity |
+-----+
| 52 | 52 | 1.0 |
| 52 | 10 | 0.8748674392700195 |
| 52 | 23 | 0.8551455140113831 |
| 52 | 26 | 0.8493421673774719 |
| 52 | 47 | 0.8411962985992432 |
| 52 | 25 | 0.8281504511833191 |
| 52 | 43 | 0.8202780485153198 |
| 52 | 24 | 0.8179885745048523 |
| 52 | 8 | 0.796689510345459 |
| 52 | 9 | 0.7947834134101868 |
| 41 | 41 | 1.0 |
| 41 | 197 | 0.9653506875038147 |
| 41 | 84 | 0.9552277326583862 |
| 41 | 157 | 0.9465565085411072 |
| 41 | 65 | 0.9287481307983398 |
| 41 | 248 | 0.9177336096763611 |
| 41 | 315 | 0.9043129086494446 |
| 41 | 92 | 0.8998928070068359 |
| 41 | 297 | 0.8897411227226257 |
| 41 | 50 | 0.8810243010520935 |
+-----+
```

### 8.4.8 Inferring a Graphlet Vector

You can infer the vector representation for a given new graphlet as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg4j> var graphlet = session.readGraphWithProperties("<path>/
<graphletConfig.json>")
```

```
opg4j> var inferredVector = model.inferGraphletVector(graphlet)
opg4j> inferredVector.print()
```

## Java

```
PgxGraph graphlet = session.readGraphWithProperties("<path>/
<graphletConfig.json>");
PgxFrame inferredVector = model.inferGraphletVector(graphlet);
inferredVector.print();
```

## Python

```
graphlet = session.read_graph_with_properties("<path>/
<graphletConfig.json>")
inferred_vector = model.infer_graphlet_vector(graphlet)
inferred_vector.print()
```

---

The schema for the `inferredVector` will be similar to the following output:

```
+-----+
| graphlet | embedding |
+-----+
```

### 8.4.9 Inferring Vectors for a Graphlet Batch

You can infer the vector representations for multiple graphlets (specified with different graph-ids in a graph) as described in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> var graphlet = session.readGraphWithProperties("<path>/
<graphletConfig.json>")
opg4j> var inferredVectorBatched =
model.inferGraphletVectorBatched(graphlets)
opg4j> inferredVectorBatched.print()
```

#### Java

```
PgxGraph graphlet = session.readGraphWithProperties("<path>/
<graphletConfig.json>");
```

```
PgxFrame inferredVectorBatched = model.inferGraphletVectorBatched(graphlets);
inferredVector.print();
```

## Python

```
graphlets = session.read_graph_with_properties("<path>/
<graphletConfig.json>")
inferred_vector_batched = model.infer_graphlet_vector_batched(graphlets)
inferred_vector_batched.print()
```

---

The schema is same as for `inferGraphletVector` but with more rows corresponding to the input graphlets.

### 8.4.10 Storing a Trained Pg2vec Model

You can store models in database. The models get stored as a row inside a model store table.

The following code shows how to store a trained Pg2vec model in database in a specific model store table:

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg4j> model.export().db().
 modelstore("modelstoretablename"). // name of the model store
table
 modelname("model"). // model name (primary key
of model store table)
 description("a model description"). // description to store
alongside the model
 store()
```

#### Java

```
model.export().db()
 .modelstore("modelstoretablename") // name of the model store table
 .modelname("model") // model name (primary key of model
store table)
 .description("a model description") // description to store alongside
the model
 .store();
```

## Python

```
model.export().db(model_store="modeltablename", # name of
the model store table
 model_name="model", # model
name (primary key of model store table)
 model_description="a model description") #
description to store alongside the model
```

 **Note:**

All the preceding examples assume that you are storing the model in the current logged in database. If you must store the model in a different database then refer to the examples in [Storing a Trained Model in Another Database](#).

### 8.4.11 Loading a Pre-Trained Pg2vec Model

You can load models from a database.

You can load a pre-trained Pg2vec model from a model store table in database as described in the following:

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg4j> var model = analyst.loadPg2vecModel().db().
modelstore("modeltablename"). // name of the model
store table
modelname("model"). // model name (primary
key of model store table)
load()
```

#### Java

```
Pg2vecModel model = analyst.loadPg2vecModel().db()
.modelstore("modeltablename") // name of the model store table
.modelname("model") // model name (primary key of model
store table)
.load();
```

## Python

```
model = analyst.get_pg2vec_model_loader(). \
 db(model_store="modelstoretablename", # name of the model
 store table
 model_name="model") # model name
(primary key of model store table)
```

 **Note:**

All the preceding examples assume that you are loading the model from the current logged in database. If you must load the model from a different database then refer to the examples in [Loading a Pre-Trained Model From Another Database](#).

## 8.4.12 Destroying a Pg2vec Model

You can destroy a Pg2vec model as described in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> model.destroy()
```

### Java

```
model.destroy();
```

### Python

```
model.destroy()
```

# 9

## OPG\_APIS Package Subprograms

The OPG\_APIS package contains subprograms (functions and procedures) for working with property graphs in an Oracle database.

To use the subprograms in this chapter, you must understand the conceptual and usage information in earlier chapters of this book.

This chapter provides reference information about the subprograms, in alphabetical order.

- [OPG\\_APIS.ANALYZE\\_PG](#)
- [OPG\\_APIS.CF](#)
- [OPG\\_APIS.CF\\_CLEANUP](#)
- [OPG\\_APIS.CF\\_PREP](#)
- [OPG\\_APIS.CLEAR\\_PG](#)
- [OPG\\_APIS.CLEAR\\_PG\\_INDICES](#)
- [OPG\\_APIS.CLONE\\_GRAPH](#)
- [OPG\\_APIS.COUNT\\_TRIANGLE](#)
- [OPG\\_APIS.COUNT\\_TRIANGLE\\_CLEANUP](#)
- [OPG\\_APIS.COUNT\\_TRIANGLE\\_PREP](#)
- [OPG\\_APIS.COUNT\\_TRIANGLE\\_RENUM](#)
- [OPG\\_APIS.CREATE\\_EDGES\\_TEXT\\_IDX](#)
- [OPG\\_APIS.CREATE\\_PG](#)
- [OPG\\_APIS.CREATE\\_PG\\_SNAPSHOT\\_TAB](#)
- [OPG\\_APIS.CREATE\\_PG\\_TEXTIDX\\_TAB](#)
- [OPG\\_APIS.CREATE\\_STAT\\_TABLE](#)
- [OPG\\_APIS.CREATE\\_SUB\\_GRAPH](#)
- [OPG\\_APIS.CREATE\\_VERTICES\\_TEXT\\_IDX](#)
- [OPG\\_APIS.DROP\\_EDGES\\_TEXT\\_IDX](#)
- [OPG\\_APIS.DROP\\_PG](#)
- [OPG\\_APIS.DROP\\_PG\\_VIEW](#)
- [OPG\\_APIS.DROP\\_VERTICES\\_TEXT\\_IDX](#)
- [OPG\\_APIS.ESTIMATE\\_TRIANGLE\\_RENUM](#)
- [OPG\\_APIS.EXP\\_EDGE\\_TAB\\_STATS](#)
- [OPG\\_APIS.EXP\\_VERTEX\\_TAB\\_STATS](#)
- [OPG\\_APIS.FIND\\_CC\\_MAPPING\\_BASED](#)
- [OPG\\_APIS.FIND\\_CLUSTERS\\_CLEANUP](#)

- OPG\_APIS.FIND\_CLUSTERS\_PREP
- OPG\_APIS.FIND\_SP
- OPG\_APIS.FIND\_SP\_CLEANUP
- OPG\_APIS.FIND\_SP\_PREP
- OPG\_APIS.GET\_BUILD\_ID
- OPG\_APIS.GET\_GEOMETRY\_FROM\_V\_COL
- OPG\_APIS.GET\_GEOMETRY\_FROM\_V\_T\_COLS
- OPG\_APIS.GET\_LATLONG\_FROM\_V\_COL
- OPG\_APIS.GET\_LATLONG\_FROM\_V\_T\_COLS
- OPG\_APIS.GET\_LONG\_LAT\_GEOMETRY
- OPG\_APIS.GET\_LATLONG\_FROM\_V\_COL
- OPG\_APIS.GET\_LONGLAT\_FROM\_V\_T\_COLS
- OPG\_APIS.GET\_SCN
- OPG\_APIS.GET\_VERSION
- OPG\_APIS.GET\_WKTGEOMETRY\_FROM\_V\_COL
- OPG\_APIS.GET\_WKTGEOMETRY\_FROM\_V\_T\_COLS
- OPG\_APIS.GRANT\_ACCESS
- OPG\_APIS.IMP\_EDGE\_TAB\_STATS
- OPG\_APIS.IMP\_VERTEX\_TAB\_STATS
- OPG\_APIS.PR
- OPG\_APIS.PR\_CLEANUP
- OPG\_APIS.PR\_PREP
- OPG\_APIS.PREPARE\_TEXT\_INDEX
- OPG\_APIS.RENAME\_PG
- OPG\_APIS.SPARSIFY\_GRAPH
- OPG\_APIS.SPARSIFY\_GRAPH\_CLEANUP
- OPG\_APIS.SPARSIFY\_GRAPH\_PREP

## 9.1 OPG\_APIS.ANALYZE\_PG

### Format

```
OPG_APIS.ANALYZE_PG(
 graph_name IN VARCHAR2,
 estimate_percent IN NUMBER,
 method_opt IN VARCHAR2,
 degree IN NUMBER,
 cascade IN BOOLEAN,
 no_invalidate IN BOOLEAN,
 force IN BOOLEAN DEFAULT FALSE,
 options IN VARCHAR2 DEFAULT NULL);
```



**Description**

Hathers, for a given property graph, statistics for the VT\$, GE\$, IT\$, and GT\$ tables.

**Parameters****graph\_name**

Name of the property graph.

**estimate\_percent**

Percentage of rows to estimate in the schema tables (NULL means compute). The valid range is [0.000001,100]. Use the constant `DBMS_STATS.AUTO_SAMPLE_SIZE` to have Oracle Database determine the appropriate sample size for good statistics. This is the usual default.

**mrthod\_opt**

Accepts either of the following options, or both in combination, for the internal property graph schema tables:

- `FOR ALL [INDEXED | HIDDEN] COLUMNS [size_clause]`
- `FOR COLUMNS [size clause] column|attribute [size_clause] [,column|attribute [size_clause]...]`

`size_clause` is defined as `size_clause := SIZE {integer | REPEAT | AUTO | SKEWONLY}`

- `integer` : Number of histogram buckets. Must be in the range [1,254].
- `REPEAT` : Collects histograms only on the columns that already have histograms.
- `AUTO` : Oracle Database determines the columns to collect histograms based on data distribution and the workload of the columns.
- `SKEWONLY` : Oracle Database determines the columns to collect histograms based on the data distribution of the columns

`column` is defined as `column := column_name | (extension)`

- `column_name` : name of a column
- `extension`: Can be either a column group in the format of `(column_name, column_name [, ...])` or an expression.

The usual default is: `FOR ALL COLUMNS SIZE AUTO`

**degree**

Degree of parallelism for the property graph schema tables. The usual default for degree is NULL, which means use the table default value specified by the `DEGREE` clause in the `CREATE TABLE` or `ALTER TABLE` statement. Use the constant `DBMS_STATS.DEFAULT_DEGREE` to specify the default value based on the initialization parameters. The `AUTO_DEGREE` value determines the degree of parallelism automatically. This is either 1 (serial execution) or `DEFAULT_DEGREE` (the system default value based on number of CPUs and initialization parameters) according to size of the object.

**cascade**

Gathers statistics on the indexes for the property graph schema tables. Use the constant `DBMS_STATS.AUTO_CASCADE` to have Oracle Database determine whether index statistics are to be collected or not. This is the usual default.

**no\_invalidate**

If `TRUE`, does not invalidate the dependent cursors. If `FALSE`, invalidates the dependent cursors immediately. If `DBMS_STATS.AUTO_INVALIDATE` (the usual default) is in effect, Oracle Database decides when to invalidate dependent cursors.

**force**

If `TRUE`, performs the operation even if one or more underlying tables are locked.

**options**

(Reserved for future use.)

**Usage Notes**

Only the owner of the property graph can call this procedure.

**Examples**

The following example gather statistics for property graph `mypg`.

```
EXECUTE OPG_APIS.ANALYZE_PG('mypg', estimate_percent=> 0.001, method_opt=>'FOR
ALL COLUMNS SIZE AUTO', degree=>4, cascade=>true, no_invalidate=>>false,
force=>true, options=>NULL);
```

## 9.2 OPG\_APIS.CF

**Format**

```
OPG_APIS.CF(
 edge_tab_name IN VARCHAR2,
 edge_label IN VARCHAR2,
 rating_property IN VARCHAR2,
 iterations IN NUMBER DEFAULT 10,
 min_error IN NUMBER DEFAULT 0.001,
 k IN NUMBER DEFAULT 5,
 learning_rate IN NUMBER DEFAULT 0.0002,
 decrease_rate IN NUMBER DEFAULT 0.95,
 regularization IN NUMBER DEFAULT 0.02,
 dop IN NUMBER DEFAULT 8,
 wt_l IN/OUT VARCHAR2,
 wt_r IN/OUT VARCHAR2,
 wt_ll IN/OUT VARCHAR2,
 wt_rl IN/OUT VARCHAR2,
 wt_i IN/OUT VARCHAR2,
 wt_ld IN/OUT VARCHAR2,
 wt_rd IN/OUT VARCHAR2,
 tablespace IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL);
```

**Description**

Runs collaborative filtering using matrix factorization on the given graph. The resulting factors of the matrix product will be stored on the left and right tables.

**Parameters****edge\_tab\_name**

Name of the property graph edge table (GE\$).

**edge\_label**

Label of the edges that hold the rating property.

**rating\_property**

(Reserved for future use: Name of the rating property.)

**iterations**

Maximum number of iterations that should be performed. Default = 10.

**min\_error**

Minimal error to reach. If at some iteration the error value is lower than this value, the procedure finishes.. Default = 0.001.

**k**

Number of features for the left and right side products. Default = 5.

**learning\_rate**

Learning rate for the gradient descent. Default = 0.0002.

**decrease\_rate**

(Reserved for future use: Decrease rate if the learning rate is too large for an effective gradient descent. Default = 0.95.)

**regularization**

An additional parameter to avoid overfitting. Default = 0.02

**dop**

Degree of parallelism. Default = 8.

**wt\_l**

Name of the working table that holds the left side of the matrix factorization.

**wt\_r**

Name of the working table that holds the right side of the matrix factorization.

**wt\_l1**

Name of the working table that holds the left side intermediate step in the gradient descent.

**wt\_r1**

Name of the working table that holds the right side intermediate step in the gradient descent.

**wt\_l**

Name of the working table that holds intermediate matrix product.

**wt\_ld**

Name of the working table that holds intermediate left side delta in gradient descent.

**wt\_rd**

Name of the working table that holds intermediate right side delta in gradient descent.

**tablespace**

Name of the tablespace to use for storing intermediate data.

**options**

Additional settings for operation. An optional string with one or more (comma-separated) of the following values:

- 'INMEMORY=T' is an option for creating the schema tables with an 'inmemory' clause.
- 'IMC\_MC\_B=T' creates the schema tables with an INMEMORY MEMCOMPRESS BASIC clause.

### Usage Notes

For information about collaborative filtering with RDF data, see [SQL-Based Property Graph Analytics](#), especially [Collaborative Filtering Overview and Examples](#).

If the working tables already exist, you can specify their names for the working table-related parameters. In this case, the algorithm can continue the progress of the previous iterations without recreating the tables.

If the working tables do not exist, or if you do not want to use existing working tables, you must first call the [OPG\\_APIS.CF\\_PREP](#) procedure, which creates the necessary working tables.

The final result of the collaborative filtering algorithm are the working tables `wt_l` and `wt_r`, which are the two factors of a matrix product. These matrix factors should be used when making predictions for collaborative filtering.

If (and only if) you have no interest in keeping the output matrix factors and the current progress of the algorithm for future use, you can call the [OPG\\_APIS.CF\\_CLEANUP](#) procedure to drop all the working tables that hold intermediate tables and the output matrix factors.

### Examples

The following example calls the [OPG\\_APIS.CF\\_PREP](#) procedure to create the working tables, and then the [OPG\\_APIS.CF](#) procedures to run collaborative filtering on the `phones` graph using the edges with the `rating` label.

```

DECLARE
 wt_l varchar2(32);
 wt_r varchar2(32);
 wt_ll varchar2(32);
 wt_rl varchar2(32);
 wt_i varchar2(32);
 wt_ld varchar2(32);
 wt_rd varchar2(32);
 edge_tab_name varchar2(32) := 'phonesge$';
 edge_label varchar2(32) := 'rating';
 rating_property varchar2(32) := '';
 iterations integer := 100;
 min_error number := 0.001;
 k integer := 5;
 learning_rate number := 0.001;
 decrease_rate number := 0.95;
 regularization number := 0.02;
 dop number := 2;
 tablespace varchar2(32) := null;
 options varchar2(32) := null;
BEGIN
 opg_apis.cf_prep(edge_tab_name,wt_l,wt_r,wt_ll,wt_rl,wt_i,wt_ld,wt_rd);
 opg_apis.cf(edge_tab_name,edge_label,rating_property,iterations,min_error,k,
 learning_rate,decrease_rate,regularization,dop,
 wt_l,wt_r,wt_ll,wt_rl,wt_i,wt_ld,wt_rd,tablespace,options);

```

```
END;
/
```

The following example assumes that OPG\_APIS.CF\_PREP had been run previously, and it specifies the various working tables that were created during that run. In this case, the preceding example automatically assigned suffixes like '\$\$CFL57' to the names of the working tables. (The output names can be printed when they are generated or be user-defined in the call to OPG\_APIS.CF\_PREP.) Thus, the following example can run more iterations of the algorithm using OPG\_APIS.CF without needing to call OPG\_APIS.CF\_PREP first, thereby continuing the progress of the previous run.

```
DECLARE
 wt_l varchar2(32) = 'phonesge$$CFL57';
 wt_r varchar2(32) = 'phonesge$$CFR57';
 wt_ll varchar2(32) = 'phonesge$$CFL157';
 wt_rl varchar2(32) = 'phonesge$$CFR157';
 wt_i varchar2(32) = 'phonesge$$CFI57';
 wt_ld varchar2(32) = 'phonesge$$CFLD57';
 wt_rd varchar2(32) = 'phonesge$$CFRD57';
 edge_tab_name varchar2(32) := 'phonesge$';
 edge_label varchar2(32) := 'rating';
 rating_property varchar2(32) := '';
 iterations integer := 100;
 min_error number := 0.001;
 k integer := 5;
 learning_rate number := 0.001;
 decrease_rate number := 0.95;
 regularization number := 0.02;
 dop number := 2;
 tablespace varchar2(32) := null;
 options varchar2(32) := null;
BEGIN
 opg_apis.cf(edge_tab_name,edge_label,rating_property,iterations,min_error,k,
 learning_rate,decrease_rate,regularization,dop,
 wt_l,wt_r,wt_ll,wt_rl,wt_i,wt_ld,wt_rd,tablespace,options);
END;
/
```

## 9.3 OPG\_APIS.CF\_CLEANUP

### Format

```
OPG_APIS.CF_CLEANUP(
 wt_l IN/OUT VARCHAR2,
 wt_r IN/OUT VARCHAR2,
 wt_ll IN/OUT VARCHAR2,
 wt_rl IN/OUT VARCHAR2,
 wt_i IN/OUT VARCHAR2,
 wt_ld IN/OUT VARCHAR2,
 wt_rd IN/OUT VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Performs cleanup work after graph collaborative filtering has been done. All the working tables that hold intermediate tables and the output matrix factors are dropped.

## Parameters

### **edge\_tab\_name**

Name of the property graph edge table (GE\$).

### **wt\_l**

Name of the working table that holds the left side of the matrix factorization.

### **wt\_r**

Name of the working table that holds the right side of the matrix factorization.

### **wt\_l1**

Name of the working table that holds the left side intermediate step in the gradient descent.

### **wt\_r1**

Name of the working table that holds the right side intermediate step in the gradient descent.

### **wt\_l**

Name of the working table that holds intermediate matrix product.

### **wt\_ld**

Name of the working table that holds intermediate left side delta in gradient descent.

### **wt\_rd**

Name of the working table that holds intermediate right side delta in gradient descent.

### **options**

(Reserved for future use.)

## Usage Notes

Call this procedure only when you have no interest in keeping the output matrix factors and the current progress of the algorithm for future use.

Do **not** call this procedure if more predictions will be made using the resulting product factors (`wt_l` and `wt_r` tables), unless you have previously made backup copies of these two tables.

See also the information about the [OPG\\_APIS.CF](#) procedure.

## Examples

The following example drops the working tables that were created in the example for the [OPG\\_APIS.CF\\_PREP](#) procedure.

```
DECLARE
 wt_l varchar2(32) = 'phonesge$$CFL57';
 wt_r varchar2(32) = 'phonesge$$CFR57';
 wt_l1 varchar2(32) = 'phonesge$$CFL157';
 wt_r1 varchar2(32) = 'phonesge$$CFR157';
 wt_i varchar2(32) = 'phonesge$$CFI57';
 wt_ld varchar2(32) = 'phonesge$$CFLD57';
 wt_rd varchar2(32) = 'phonesge$$CFRD57';
BEGIN
 opg_apis.cf_cleanup('phonesge$',wt_l,wt_r,wt_l1,wt_r1,wt_i,wt_ld,wt_rd);
```

```
END;
/
```

## 9.4 OPG\_APIS.CF\_PREP

### Format

```
OPG_APIS.CF_PREP(
 wt_l IN/OUT VARCHAR2.
 wt_r IN/OUT VARCHAR2.
 wt_l1 IN/OUT VARCHAR2.
 wt_r1 IN/OUT VARCHAR2.
 wt_i IN/OUT VARCHAR2.
 wt_ld IN/OUT VARCHAR2.
 wt_rd IN/OUT VARCHAR2.
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Performs preparation work, including creating the necessary intermediate tables, for a later call to the [OPG\\_APIS.CF](#) procedure that will perform collaborative filtering.

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table (GE\$).

#### **wt\_l**

Name of the working table that holds the left side of the matrix factorization.

#### **wt\_r**

Name of the working table that holds the right side of the matrix factorization.

#### **wt\_l1**

Name of the working table that holds the left side intermediate step in the gradient descent.

#### **wt\_r1**

Name of the working table that holds the right side intermediate step in the gradient descent.

#### **wt\_i**

Name of the working table that holds intermediate matrix product.

#### **wt\_ld**

Name of the working table that holds intermediate left side delta in gradient descent.

#### **wt\_rd**

Name of the working table that holds intermediate right side delta in gradient descent.

#### **options**

Additional settings for operation. An optional string with one or more (comma-separated) of the following values:

- 'INMEMORY=T' is an option for creating the schema tables with an 'inmemory' clause.
- 'IMC\_MC\_B=T' creates the schema tables with an INMEMORY MEMCOMPRESS BASIC clause.

## Usage Notes

The names of the working tables can be specified or left as null parameters. If the name of any working table parameter is not specified, a name is automatically generated and is returned as an OUT parameter. The working table names can be used when you call the [OPG\\_APIS.CF](#) procedure to run the collaborative filtering algorithm.

See also the Usage Notes and Examples for [OPG\\_APIS.CF](#).

## Examples

The following example creates the working tables for a graph named `phones`, and it prints the names that were automatically generated for the working tables.

```
DECLARE
 wt_l varchar2(32);
 wt_r varchar2(32);
 wt_ll varchar2(32);
 wt_rl varchar2(32);
 wt_i varchar2(32);
 wt_ld varchar2(32);
 wt_rd varchar2(32);
BEGIN
 opg_apis.cf_prep('phonesge$',wt_l,wt_r,wt_ll,wt_rl,wt_i,wt_ld,wt_rd);
 dbms_output.put_line(' wt_l ' || wt_l);
 dbms_output.put_line(' wt_r ' || wt_r);
 dbms_output.put_line(' wt_ll ' || wt_ll);
 dbms_output.put_line(' wt_rl ' || wt_rl);
 dbms_output.put_line(' wt_i ' || wt_i);
 dbms_output.put_line(' wt_ld ' || wt_ld);
 dbms_output.put_line(' wt_rd ' || wt_rd);
END;
/
```

## 9.5 OPG\_APIS.CLEAR\_PG

### Format

```
OPG_APIS.CLEAR_PG(
 graph_name IN VARCHAR2);
```

### Description

Clears all data from a property graph.

### Parameters

#### **graph\_name**

Name of the property graph.

### Usage Notes

This procedure removes all data in the property graph by deleting data in the graph tables (VT\$, GE\$, and so on).



### Examples

The following example removes all data from the property graph named `mypg`.

```
EXECUTE OPG_APIS.CLEAR_PG('mypg');
```

## 9.6 OPG\_APIS.CLEAR\_PG\_INDICES

### Format

```
OPG_APIS.CLEAR_PG(
 graph_name IN VARCHAR2);
```

### Description

Removes all text index metadata in the IT\$ table of the property graph.

### Parameters

#### **graph\_name**

Name of the property graph.

### Usage Notes

This procedure does not actually remove text index data

### Examples

The following example removes all index metadata of the property graph named `mypg`.

```
EXECUTE OPG_APIS.CLEAR_PG_INDICES('mypg');
```

## 9.7 OPG\_APIS.CLONE\_GRAPH

### Format

```
OPG_APIS.CLONE_GRAPH(
 orgGraph IN VARCHAR2,
 newGraph IN VARCHAR2,
 dop IN INTEGER DEFAULT 4,
 num_hash_ptns IN INTEGER DEFAULT 8,
 tbs IN VARCHAR2 DEFAULT NULL);
```

### Description

Makes a clone of the original graph, giving the new graph a new name.

### Parameters

#### **orgGraph**

Name of the original property graph.

#### **newGraph**

Name of the new (clone) property graph.

**dop**

Degree of parallelism for the operation.

**num\_hash\_ptns**

Number of hash partitions used to partition the vertices and edges tables. It is recommended to use a power of 2 (2, 4, 8, 16, and so on).

**tbs**

Name of the tablespace to hold all the graph data and index data.

**Usage Notes**

The original property graph must already exist in the database.

**Examples**

The following example creates a clone graph named `mypgclone` from the property graph `mypg` in the tablespace `my_ts` using a degree of parallelism of 4 and 8 partitions.

```
EXECUTE OPG_APIS.CLONE_GRAPH('mypg', 'mypgclone', 4, 8, 'my_ts');
```

## 9.8 OPG\_APIS.COUNT\_TRIANGLE

**Format**

```
OPG_APIS.COUNT_TRIANGLE(
 edge_tab_name IN VARCHAR2,
 wt_und IN OUT VARCHAR2,
 num_sub_ptns IN NUMBER DEFAULT 1,
 dop IN INTEGER DEFAULT 1,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL
) RETURN NUMBER;
```

**Description**

Performs triangle counting in property graph.

**Parameters****edge\_tab\_name**

Name of the property graph edge table.

**wt\_und**

A working table holding an undirected version of the graph.

**num\_sub\_ptns**

Number of logical subpartitions used in calculating triangles . Must be a positive integer, power of 2 (1, 2, 4, 8, ...). For a graph with a relatively small maximum degree, use the value 1 (the default).

**dop**

Degree of parallelism for the operation. The default is 1.

**tbs**

Name of the tablespace to hold the data stored in working tables.

**options**

Additional settings for the operation:

- 'PDML=T' enables parallel DML.

**Usage Notes**

The property graph edge table must exist in the database, and the [OPG\\_APIS.COUNT\\_TRIANGLE\\_PREP](#) procedure must already have been executed.

**Examples**

The following example performs triangle counting in the property graph named `connections`

```
set serveroutput on
DECLARE
 wt1 varchar2(100); -- intermediate working table
 wt2 varchar2(100);
 wt3 varchar2(100);
 n number;
BEGIN
 opg_apis.count_triangle_prep('connectionsGE$', wt1, wt2, wt3);
 n := opg_apis.count_triangle(
 'connectionsGE$',
 wt1,
 num_sub_ptns=>1,
 dop=>2,
 tbs => 'MYPG_TS',
 options=>'PDML=T'
);
 dbms_output.put_line('total number of triangles ' || n);
END;
/
```

## 9.9 OPG\_APIS.COUNT\_TRIANGLE\_CLEANUP

**Format**

```
COUNT_TRIANGLE_CLEANUP(
 edge_tab_name IN VARCHAR2,
 wt_undBM IN VARCHAR2,
 wt_rnmap IN VARCHAR2,
 wt_undAM IN VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

**Description**

Cleans up and drops the temporary working tables used for triangle counting.

**Parameters****edge\_tab\_name**

Name of the property graph edge table.

**wt\_undBM**

A working table holding an undirected version of the original graph (before renumbering optimization).

**wt\_rnmap**

A working table that is a mapping table for renumbering optimization.

**wt\_undAM**

A working table holding the undirected version of the graph data after applying the renumbering optimization.

**options**

Additional settings for operation. An optional string with one or more (comma-separated) of the following values:

- PDML=T enables parallel DML.

**Usage Notes**

You should use this procedure to clean up after triangle counting.

The working tables must exist in the database.

**Examples**

The following example performs triangle counting in the property graph named `connections`, and drops the working table after it has finished.

```
set serveroutput on

DECLARE
 wt1 varchar2(100); -- intermediate working table
 wt2 varchar2(100);
 wt3 varchar2(100);
 n number;
BEGIN
 opg_apis.count_triangle_prep('connectionsGE$', wt1, wt2, wt3);
 n := opg_apis.count_triangle_renum(
 'connectionsGE$',
 wt1,
 wt2,
 wt3,
 num_sub_ptns=>1,
 dop=>2,
 tbs => 'MYPG_TS',
 options=>'PDML=T'
);
 dbms_output.put_line('total number of triangles ' || n);
 opg_apis.count_triangle_cleanup('connectionsGE$', wt1, wt2, wt3);
END;
/
```

## 9.10 OPG\_APIS.COUNT\_TRIANGLE\_PREP

**Format**

```
OPG_APIS.COUNT_TRIANGLE_PREP(
 edge_tab_name IN VARCHAR2,
 wt_undBM IN OUT VARCHAR2,
 wt_rnmap IN OUT VARCHAR2,
 wt_undAM IN OUT VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

**Description**

Prepares for running triangle counting.

**Parameters****edge\_tab\_name**

Name of the property graph edge table.

**wt\_undBM**

A working table holding an undirected version of the original graph (before renumbering optimization).

**wt\_rnmap**

A working table that is a mapping table for renumbering optimization.

**wt\_undAM**

A working table holding the undirected version of the graph data after applying the renumbering optimization.

**options**

Additional settings for operation. An optional string with one or more (comma-separated) of the following values:

- CREATE\_UNDIRECTED=T
- REUSE\_UNDIRECTED\_TAB=T

**Usage Notes**

The property graph edge table must exist in the database.

**Examples**

The following example prepares for triangle counting in a property graph named connections.

```
set serveroutput on

DECLARE
 wt1 varchar2(100); -- intermediate working table
 wt2 varchar2(100);
 wt3 varchar2(100);
 n number;
BEGIN
 opg_apis.count_triangle_prep('connectionsGE$', wt1, wt2, wt3);

 n := opg_apis.count_triangle_renum(
 'connectionsGE$',
 wt1,
 wt2,
 wt3,
 num_sub_ptns=>1,
 dop=>2,
 tbs => 'MYPG_TS',
 options=>'CREATE_UNDIRECTED=T,REUSE_UNDIREC_TAB=T'
);
 dbms_output.put_line('total number of triangles ' || n);
```

```
END;
/
```

## 9.11 OPG\_APIS.COUNT\_TRIANGLE\_RENUM

### Format

```
COUNT_TRIANGLE_RENUM(
 edge_tab_name IN VARCHAR2,
 wt_undBM IN VARCHAR2,
 wt_rnmap IN VARCHAR2,
 wt_undAM IN VARCHAR2,
 num_sub_ptns IN INTEGER DEFAULT 1,
 dop IN INTEGER DEFAULT 1,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL
) RETURN NUMBER;
```

### Description

Performs triangle counting in property graph, with the optimization of renumbering the vertices of the graph by their degree.

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table.

#### **wt\_undBM**

A working table holding an undirected version of the original graph (before renumbering optimization).

#### **wt\_rnmap**

A working table that is a mapping table for renumbering optimization.

#### **wt\_undAM**

A working table holding the undirected version of the graph data after applying the renumbering optimization.

#### **num\_sub\_ptns**

Number of logical subpartitions used in calculating triangles . Must be a positive integer, power of 2 (1, 2, 4, 8, ...). For a graph with a relatively small maximum degree, use the value 1 (the default).

#### **dop**

Degree of parallelism for the operation. The default is 1 (no parallelism).

#### **tbs**

Name of the tablespace to hold the data stored in working tables.

#### **options**

Additional settings for operation. An optional string with one or more (comma-separated) of the following values:

- PDML=T enables parallel DML.

### Usage Notes

This function makes the algorithm run faster, but requires more space.

The property graph edge table must exist in the database, and the [OPG\\_APIS.COUNT\\_TRIANGLE\\_PREP](#) procedure must already have been executed.

### Examples

The following example performs triangle counting in the property graph named `connections`. It does not perform the cleanup after it finishes, so you can count triangles again on the same graph without calling the preparation procedure.

```

set serveroutput on

DECLARE
 wt1 varchar2(100); -- intermediate working table
 wt2 varchar2(100);
 wt3 varchar2(100);
 n number;
BEGIN
 opg_apis.count_triangle_prep('connectionsGE$', wt1, wt2, wt3);
 n := opg_apis.count_triangle_renum(
 'connectionsGE$',
 wt1,
 wt2,
 wt3,
 num_sub_ptns=>1,
 dop=>2,
 tbs => 'MYPG_TS',
 options=>'PDML=T'
);
 dbms_output.put_line('total number of triangles ' || n);
END;
/

```

## 9.12 OPG\_APIS.CREATE\_EDGES\_TEXT\_IDX

### Format

```

OPG_APIS.CREATE_EDGES_TEXT_IDX(
 graph_owner IN VARCHAR2,
 graph_name IN VARCHAR2,
 pref_owner IN VARCHAR2 DEFAULT NULL,
 datastore IN VARCHAR2 DEFAULT NULL,
 filter IN VARCHAR2 DEFAULT NULL,
 storage IN VARCHAR2 DEFAULT NULL,
 wordlist IN VARCHAR2 DEFAULT NULL,
 stoplist IN VARCHAR2 DEFAULT NULL,
 lexer IN VARCHAR2 DEFAULT NULL,
 dop IN INTEGER DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL,);

```

### Description

Creates a text index on a property graph edge table.

**Parameters**

**graph\_owner**  
Owner of the property graph.

**graph\_name**  
Name of the property graph.

**pref\_owner**  
Owner of the preference.

**datastore**  
The way that documents are stored.

**filter**  
The way that documents can be converted to plain text.

**storage**  
The way that the index data is stored.

**wordlist**  
The way that stem and fuzzy queries should be expanded

**stoplist**  
The words or themes that are not to be indexed.

**lexer**  
The language used for indexing.

**dop**  
The degree of parallelism used for index creation.

**options**  
Additional settings for index creation.

**Usage Notes**

The property graph must exist in the database.

You must have the ALTER SESSION privilege to run this procedure.

**Examples**

The following example creates a text index on the edge table of property graph `mypg`, which is owned by user SCOTT, using the lexer `OPG_AUTO_LEXER` and a degree of parallelism of 4.

```
EXECUTE OPG_APIS.CREATE_EDGES_TEXT_IDX('SCOTT', 'mypg', 'MDSYS', null, null,
null, null, null, 'OPG_AUTO_LEXER', 4, null);
```

## 9.13 OPG\_APIS.CREATE\_PG

**Format**

```
OPG_APIS.CREATE_PG(
 graph_name IN VARCHAR2,
 dop IN INTEGER DEFAULT NULL,
```



```
num_hash_ptns IN INTEGER DEFAULT 8,
tbs IN VARCHAR2 DEFAULT NULL,
options IN VARCHAR2 DEFAULT NULL);
```

### Description

Creates, for a given property graph name, the necessary property graph schema tables that are necessary to store data about vertices, edges, text indexes, and snapshots.

### Parameters

#### **graph\_name**

Name of the property graph.

#### **dop**

Degree of parallelism for the operation.

#### **num\_hash\_ptns**

Number of hash partitions used to partition the vertices and edges tables. It is recommended to use a power of 2 (2, 4, 8, 16, and so on).

#### **tbs**

Name of the tablespace to hold all the graph data and index data.

#### **options**

Options that can be used to customize the creation of indexes on schema tables. (One or more, comma separated.)

- 'SKIP\_INDEX=T' skips the default index creation.
- 'SKIP\_ERROR=T' ignores errors encountered during table/index creation.
- 'INMEMORY=T' creates the schema tables with an INMEMORY clause.
- 'IMC\_MC\_B=T' creates the schema tables with an INMEMORY BASIC clause.

### Usage Notes

You must have the CREATE TABLE and CREATE INDEX privileges to call this procedure.

By default, all the schema tables will be created with basic compression enabled.

### Examples

The following example creates a property graph named `mypg` in the tablespace `my_ts` using eight partitions.

```
EXECUTE OPG_APIS.CREATE_PG('mypg', 4, 8, 'my_ts');
```

## 9.14 OPG\_APIS.CREATE\_PG\_SNAPSHOT\_TAB

### Format

```
OPG_APIS.CREATE_PG_SNAPSHOT_TAB(
 graph_owner IN VARCHAR2,
 graph_name IN VARCHAR2,
 dop IN INTEGER DEFAULT NULL,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL);
```

or

```
OPG_APIS.CREATE_PG_SNAPSHOT_TAB (
 graph_name IN VARCHAR2,
 dop IN INTEGER DEFAULT NULL,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Creates, for a given property graph name, the necessary property graph schema table (<graph\_name>SS\$) that stores data about snapshots for the graph.

### Parameters

#### graph\_owner

Name of the owner of the property graph.

#### graph\_name

Name of the property graph.

#### dop

Degree of parallelism for the operation.

#### tbs

Name of the tablespace to hold all the graph snapshot data and associated index.

#### options

Additional settings for the operation:

- 'INMEMORY=T' is an option for creating the schema tables with an 'inmemory' clause.
- 'IMC\_MC\_B=T' creates the schema tables with an INMEMORY MEMCOMPRESS BASIC clause.

### Usage Notes

You must have the CREATE TABLE privilege to call this procedure.

The created snapshot table has the following structure, which may change between releases.

Name	Null?	Type
SSID	NOT NULL	NUMBER
CONTENTS		BLOB
SS_FILE		BINARY FILE LOB
TS		TIMESTAMP(6) WITH TIME ZONE
SS_COMMENT		VARCHAR2(512)

By default, all schema tables will be created with basic compression enabled.

### Examples

The following example creates a snapshot table for property graph `mypg` in the current schema, with a degree of parallelism of 4 and using the `MY_TS` tablespace.

```
EXECUTE OPG_APIS.CREATE_PG_SNAPSHOT_TAB('mypg', 4, 'my_ts');
```

## 9.15 OPG\_APIS.CREATE\_PG\_TEXTIDX\_TAB

### Format

```
OPG_APIS.CREATE_PG_TEXTIDX_TAB(
 graph_owner IN VARCHAR2,
 graph_name IN VARCHAR2,
 dop IN INTEGER DEFAULT NULL,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL);
```

or

```
OPG_APIS.CREATE_PG_TEXTIDX_TAB(
 graph_name IN VARCHAR2,
 dop IN INTEGER DEFAULT NULL,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Creates, for a given property graph name, the necessary property graph text index schema table (<graph\_name>IT\$) that stores data for managing text index metadata for the graph.

### Parameters

#### **graph\_owner**

Name of the owner of the property graph.

#### **graph\_name**

Name of the property graph.

#### **dop**

Degree of parallelism for the operation.

#### **tbs**

Name of the tablespace to hold all the graph index metadata and associated index.

#### **options**

Additional settings for the operation:

- 'INMEMORY=T' is an option for creating the schema tables with an 'inmemory' clause.
- 'IMC\_MC\_B=T' creates the schema tables with an INMEMORY MEMCOMPRESS BASIC clause.

### Usage Notes

You must have the CREATE TABLE privilege to call this procedure.

The created index metadata table has the following structure, which may change between releases.

```
(
 EIN nvarchar2(80) not null, -- index name
 ET number, -- entity type 1 - vertex, 2 -edge
 IT number, -- index type 1 - auto 0 - manual
 SE number, -- search engine 1 -solr, 0 - lucene
 K nvarchar2(3100), -- property key use an empty space when
```

```

there is no K/V
 DT number, -- directory type 1 - MMAP, 2 -
FS, 3 - JDBC
 LOC nvarchar2(3100), -- directory location (1, 2)
 NUMDIRS number, -- property key used to index CAN
BE NULL
 VERSION nvarchar2(100), -- lucene version
 USEDT number, -- user data type (1 or 0)
 STOREF number, -- store fields into lucene
 CF nvarchar2(3100), -- configuration name
 SS nvarchar2(3100), -- solr server url
 SA nvarchar2(3100), -- solr server admin url
 ZT number, -- zookeeper timeout
 SH number, -- number of shards
 RF number, -- replication factor
 MS number, -- maximum shards per node
 PO nvarchar2(3100), -- preferred owner oracle text
 DS nvarchar2(3100), -- datastore
 FIL nvarchar2(3100), -- filter
 STR nvarchar2(3100), -- storage
 WL nvarchar2(3100), -- word list
 SL nvarchar2(3100), -- stop list
 LXR nvarchar2(3100), -- lexer
 OPTS nvarchar2(3100), -- options
 primary key (EIN, K, ET)
)

```

By default, all schema tables will be created with basic compression enabled.

### Examples

The following example creates a property graph text index metadata table for property graph `mypg` in the current schema, with a degree of parallelism of 4 and using the `MY_TS` tablespace.

```
EXECUTE OPG_APIS.CREATE_PG_TEXTIDX_TAB('mypg', 4, 'my_ts');
```

## 9.16 OPG\_APIS.CREATE\_STAT\_TABLE

### Format

```
OPG_APIS.CREATE_STAT_TABLE(
 stattab IN VARCHAR2,
 tblspace IN VARCHAR2 DEFAULT NULL);
```

### Description

Creates a table that can hold property graph statistics.

### Parameters

#### **stattab**

Name of the table to hold statistics

#### **tblspace**

Name of the tablespace to hold the statistics table. If none is specified, then the statistics table will be created in the user's default tablespace.

## Usage Notes

You must have the CREATE TABLE privilege to call this procedure.

The statistics table has the following columns. Note that the columns and their types may vary between releases.

Name	Null?	Type
STATID		VARCHAR2 (128)
TYPE		CHAR (1)
VERSION		NUMBER
FLAGS		NUMBER
C1		VARCHAR2 (128)
C2		VARCHAR2 (128)
C3		VARCHAR2 (128)
C4		VARCHAR2 (128)
C5		VARCHAR2 (128)
C6		VARCHAR2 (128)
N1		NUMBER
N2		NUMBER
N3		NUMBER
N4		NUMBER
N5		NUMBER
N6		NUMBER
N7		NUMBER
N8		NUMBER
N9		NUMBER
N10		NUMBER
N11		NUMBER
N12		NUMBER
N13		NUMBER
D1		DATE
T1		TIMESTAMP (6) WITH TIME ZONE
R1		RAW (1000)
R2		RAW (1000)
R3		RAW (1000)
CH1		VARCHAR2 (1000)
CL1		CLOB

## Examples

The following example creates a statistics table named `mystat`.

```
EXECUTE OPG_APIS.CREATE_STAT_TABLE('mystat', null);
```

## 9.17 OPG\_APIS.CREATE\_SUB\_GRAPH

### Format

```
OPG_APIS.CREATE_SUB_GRAPH (
 graph_owner IN VARCHAR2,
 orgGraph IN VARCHAR2,
 newGraph IN VARCHAR2,
 nSrc IN NUMBER,
 depth IN NUMBER);
```

**Description**

Creates a subgraph, which is an expansion from a given vertex. The depth of expansion is customizable.

**Parameters****graph\_owner**

Owner of the property graph.

**orgGraph**

Name of the original property graph.

**newGraph**

Name of the subgraph to be created from the original graph.

**nSrc**

Vertex ID: the subgraph will be created by expansion from this vertex. For example, `nSrc = 1` starts the expansion from the vertex with ID 1.

**depth**

Depth of expansion: the expansion, following outgoing edges, will include all vertices that are within `depth` hops away from vertex `nSrc`. For example, `depth = 2` causes the to should include all vertices that are within 2 hops away from vertex `nSrc` (vertex ID 1 in the preceding example).

**Usage Notes**

The original property graph must exist in the database.

**Examples**

The following example creates a subgraph `mypgsub` from the property graph `mypg` whose owner is SCOTT. The subgraph includes vertex 1 and all vertices that are reachable from the vertex with ID 1 in 2 hops.

```
EXECUTE OPG_APIS.CREATE_SUB_GRAPH('SCOTT', 'mypg', 'mypgsub', 1, 2);
```

## 9.18 OPG\_APIS.CREATE\_VERTICES\_TEXT\_IDX

**Format**

```
OPG_APIS.CREATE_VERTICES_TEXT_IDX(
 graph_owner IN VARCHAR2,
 graph_name IN VARCHAR2,
 pref_owner IN VARCHAR2 DEFAULT NULL,
 datastore IN VARCHAR2 DEFAULT NULL,
 filter IN VARCHAR2 DEFAULT NULL,
 storage IN VARCHAR2 DEFAULT NULL,
 wordlist IN VARCHAR2 DEFAULT NULL,
 stoplist IN VARCHAR2 DEFAULT NULL,
 lexer IN VARCHAR2 DEFAULT NULL,
 dop IN INTEGER DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL,);
```

**Description**

Creates a text index on a property graph vertex table.

**Parameters****graph\_owner**

Owner of the property graph.

**graph\_name**

Name of the property graph.

**pref\_owner**

Owner of the preference.

**datastore**

The way that documents are stored.

**filter**

The way that documents can be converted to plain text.

**storage**

The way that the index data is stored.

**wordlist**

The way that stem and fuzzy queries should be expanded

**stoplist**

The words or themes that are not to be indexed.

**lexer**

The language used for indexing.

**dop**

The degree of parallelism used for index creation.

**options**

Additional settings for index creation.

**Usage Notes**

The original property graph must exist in the database.

You must have the ALTER SESSION privilege to run this procedure.

**Examples**

The following example creates a text index on the vertex table of property graph `mypg`, which is owned by user `SCOTT`, using the lexer `OPG_AUTO_LEXER` and a degree of parallelism of 4.

```
EXECUTE OPG_APIS.CREATE_VERTICES_TEXT_IDX('SCOTT', 'mypg', null, null, null, null,
null, null, 'OPG_AUTO_LEXER', 4, null);
```

## 9.19 OPG\_APIS.DROP\_EDGES\_TEXT\_IDX

### Format

```
OPG_APIS.DROP_EDGES_TEXT_IDX(
 graph_owner IN VARCHAR2,
 graph_name IN VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Drops a text index on a property graph edge table.

### Parameters

#### **graph\_owner**

Owner of the property graph.

#### **graph\_name**

Name of the property graph.

#### **options**

Additional settings for the operation.

### Usage Notes

A text index must already exist on the property graph edge table.

### Examples

The following example drops the text index on the edge table of property graph `mypg` that is owned by user `SCOTT`.

```
EXECUTE OPG_APIS.DROP_EDGES_TEXT_IDX('SCOTT', 'mypg', null);
```

## 9.20 OPG\_APIS.DROP\_PG

### Format

```
OPG_APIS.DROP_PG(
 graph_name IN VARCHAR2);
```

### Description

Drops (deletes) a property graph.

### Parameters

#### **graph\_name**

Name of the property graph.

### Usage Notes

All the graph tables (VT\$, GE\$, and so on) will be dropped from the database.



## Examples

The following example drops the property graph named `mypg`.

```
EXECUTE OPG_APIS.DROP_PG('mypg');
```

## 9.21 OPG\_APIS.DROP\_PG\_VIEW

### Format

```
OPG_APIS.DROP_PG_VIEW(
 graph_name IN VARCHAR2);
 options IN VARCHAR2);
```

### Description

Drops (deletes) the view definition of a property graph.

### Parameters

#### **graph\_name**

Name of the property graph.

#### **options**

(Reserved for future use.)

### Usage Notes

Oracle supports creating physical property graphs and property graph views. For example, given an RDF model, it supports creating property graph views over the RDF model, so that you can run property graph analytics on top of the RDF graph.

This procedure cannot be undone.

### Examples

The following example drops the view definition of the property graph named `mypg`.

```
EXECUTE OPG_APIS.DROP_PG_VIEW('mypg');
```

## 9.22 OPG\_APIS.DROP\_VERTICES\_TEXT\_IDX

### Format

```
OPG_APIS.DROP_VERTICES_TEXT_IDX(
 graph_owner IN VARCHAR2,
 graph_name IN VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Drops a text index on a property graph vertex table.

**Parameters****graph\_owner**

Owner of the property graph.

**graph\_name**

Name of the property graph.

**options**

Additional settings for the operation.

**Usage Notes**

A text index must already exist on the property graph vertex table.

**Examples**

The following example drops the text index on the vertex table of property graph `mypg` that is owned by user `SCOTT`.

```
EXECUTE OPG_APIS.DROP_VERTICES_TEXT_IDX('SCOTT', 'mypg', null);
```

## 9.23 OPG\_APIS.ESTIMATE\_TRIANGLE\_RENUM

**Format**

```
COUNT_TRIANGLE_ESTIMATE(
 edge_tab_name IN VARCHAR2,
 wt_undBM IN VARCHAR2,
 wt_rnmap IN VARCHAR2,
 wt_undAM IN VARCHAR2,
 num_sub_ptns IN INTEGER DEFAULT 1,
 chunk_id IN INTEGER DEFAULT 1,
 dop IN INTEGER DEFAULT 1,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL
) RETURN NUMBER;
```

**Description**

Estimates the number of triangles in a property graph.

**Parameters****edge\_tab\_name**

Name of the property graph edge table.

**wt\_undBM**

A working table holding an undirected version of the original graph (before renumbering optimization).

**wt\_rnmap**

A working table that is a mapping table for renumbering optimization.

**wt\_undAM**

A working table holding the undirected version of the graph data after applying the renumbering optimization.

**num\_sub\_ptns**

Number of logical subpartitions used in calculating triangles. Must be a positive integer, power of 2 (1, 2, 4, 8, ...). For a graph with a relatively small maximum degree, use the value 1 (the default).

**chunk\_id**

The logical subpartition to be used in triangle estimation (Only this partition will be counted). It must be an integer between 0 and  $\text{num\_sub\_ptns} * \text{num\_sub\_ptns} - 1$ .

**dop**

Degree of parallelism for the operation. The default is 1 (no parallelism).

**tbs**

Name of the tablespace to hold the data stored in working tables.

**options**

Additional settings for operation. An optional string with one or more (comma-separated) of the following values:

- PDML=T enables parallel DML.

**Usage Notes**

This function counts the total triangles in a portion of size  $1 / (\text{num\_sub\_ptns} * \text{num\_sub\_ptns})$  of the graph; so to estimate the total number of triangles in the graph, you can multiply the result by  $\text{num\_sub\_ptns} * \text{num\_sub\_ptns}$ .

The property graph edge table must exist in the database, and the [OPG\\_APIS.COUNT\\_TRIANGLE\\_PREP](#) procedure must already have been executed.

**Examples**

The following example estimates the number of triangle in the property graph named `connections`. It does not perform the cleanup after it finishes, so you can count triangles again on the same graph without calling the preparation procedure.

```
set serveroutput on

DECLARE
 wt1 varchar2(100); -- intermediate working table
 wt2 varchar2(100);
 wt3 varchar2(100);
 n number;
BEGIN
 opg_apis.count_triangle_prep('connectionsGE$', wt1, wt2, wt3);
 n := opg_apis.estimate_triangle_renum(
 'connectionsGE$',
 wt1,
 wt2,
 wt3,
 num_sub_ptns=>64,
 chunk_id=>2048,
 dop=>2,
 tbs => 'MYPG_TS',
 options=>'PDML=T'
```

```

);
 dbms_output.put_line('estimated number of triangles ' || (n * 64 * 64));
END;
/

```

## 9.24 OPG\_APIS.EXP\_EDGE\_TAB\_STATS

### Format

```

OPG_APIS.EXP_EDGE_TAB_STATS (
 graph_name IN VARCHAR2,
 stattab IN VARCHAR2,
 statid IN VARCHAR2 DEFAULT NULL,
 cascade IN BOOLEAN DEFAULT TRUE,
 statown IN VARCHAR2 DEFAULT NULL,
 stat_category IN VARCHAR2 DEFAULT 'OBJECT_STATS');

```

### Description

Retrieves statistics for the edge table of a given property graph and stores them in the user-created statistics table.

### Parameters

#### **graph\_name**

Name of the property graph.

#### **stattab**

Name of the statistics table.

#### **statid**

Optional identifier to associate with these statistics within `stattab`.

#### **cascade**

If `TRUE`, column and index statistics are exported.

#### **statown**

Schema containing `stattab`.

#### **stat\_category**

Specifies what statistics to export, using a comma to separate values. The supported values are 'OBJECT\_STATS' (the default: table statistics, column statistics, and index statistics) and 'SYNOPSIS' (auxiliary statistics created when statistics are incrementally maintained).

### Usage Notes

(None.)

### Examples

The following example creates a statistics table, exports into this table the property graph edge table statistics, and issues a query to count the relevant rows for the newly created statistics.

```

EXECUTE OPG_APIS.CREATE_STAT_TABLE('mystat',null);

EXECUTE OPG_APIS.EXP_EDGE_TAB_STATS('mypg', 'mystat', 'edge_stats_id_1', true,

```

```

null, 'OBJECT_STATS');

SELECT count(1) FROM mystat WHERE statid='EDGE_STATS_ID_1';

153

```

## 9.25 OPG\_APIS.EXP\_VERTEX\_TAB\_STATS

### Format

```

OPG_APIS.EXP_VERTEX_TAB_STATS(
 graph_name IN VARCHAR2,
 stattab IN VARCHAR2,
 statid IN VARCHAR2 DEFAULT NULL,
 cascade IN BOOLEAN DEFAULT TRUE,
 statown IN VARCHAR2 DEFAULT NULL,
 stat_category IN VARCHAR2 DEFAULT 'OBJECT_STATS');

```

### Description

Retrieves statistics for the vertex table of a given property graph and stores them in the user-created statistics table.

### Parameters

#### **graph\_name**

Name of the property graph.

#### **stattab**

Name of the statistics table.

#### **statid**

Optional identifier to associate with these statistics within `stattab`.

#### **cascade**

If `TRUE`, column and index statistics are exported.

#### **statown**

Schema containing `stattab`.

#### **stat\_category**

Specifies what statistics to export, using a comma to separate values. The supported values are `'OBJECT_STATS'` (the default: table statistics, column statistics, and index statistics) and `'SYNOPSIS'` (auxiliary statistics created when statistics are incrementally maintained).

### Usage Notes

(None.)

### Examples

The following example creates a statistics table, exports into this table the property graph vertex table statistics, and issues a query to count the relevant rows for the newly created statistics.

```

EXECUTE OPG_APIS.CREATE_STAT_TABLE('mystat',null);

EXECUTE OPG_APIS.EXP_VERTEX_TAB_STATS('mypg', 'mystat', 'vertex_stats_id_1', true,

```

```
null, 'OBJECT_STATS');

SELECT count(1) FROM mystat WHERE statid='VERTEX_STATS_ID_1';

108
```

## 9.26 OPG\_APIS.FIND\_CC\_MAPPING\_BASED

### Format

```
OPG_APIS.FIND_CC_MAPPING_BASED(
 edge_tab_name IN VARCHAR2,
 wt_clusters IN OUT VARCHAR2,
 wt_undir IN OUT VARCHAR2,
 wt_cluas IN OUT VARCHAR2,
 wt_newas IN OUT VARCHAR2,
 wt_delta IN OUT VARCHAR2,
 dop IN INTEGER DEFAULT 4,
 rounds IN INTEGER DEFAULT 0,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Finds connected components in a property graph. All connected components will be stored in the `wt_clusters` table. The original graph is treated as undirected.

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table.

#### **wt\_clusters**

A working table holding the final vertex cluster mappings. This table has two columns (VID NUMBER, CLUSTER\_ID NUMBER). Column VID stores the vertex ID values, and column CLUSTER\_ID stores the corresponding cluster ID values. Cluster ID values are long integers that can have gaps between them.

If an empty name is specified, a new table will be generated, and its name will be returned.

#### **wt\_undir**

A working table holding an undirected version of the graph.

#### **wt\_cluas**

A working table holding current cluster assignments.

#### **wt\_newas**

A working table holding updated cluster assignments.

#### **wt\_delta**

A working table holding changes ("delta") in cluster assignments.

#### **dop**

Degree of parallelism for the operation. The default is 4.

**rounds**

Maximum number of iterations to perform in searching for connected components. The default value of 0 (zero) means that computation will continue until all connected components are found.

**tbs**

Name of the tablespace to hold the data stored in working tables.

**options**

Additional settings for the operation.

- 'PDML=T' enables parallel DML.

**Usage Notes**

The property graph edge table must exist in the database, and the [OPG\\_APIS.FIND\\_CLUSTERS\\_PREP](#) procedure must already have been executed.

**Examples**

The following example finds the connected components in a property graph named `mypg`.

```
DECLARE
 wtClusters varchar2(200) := 'mypg_clusters';
 wtUnDir varchar2(200);
 wtCluas varchar2(200);
 wtNewas varchar2(200);
 wtDelta varchar2(200);
BEGIN
 opg_apis.find_clusters_prep('mypgGE$', wtClusters, wtUnDir,
 wtCluas, wtNewas, wtDelta, '');
 dbms_output.put_line('working tables names ' || wtClusters || ' '
 || wtUnDir || ' ' || wtCluas || ' ' || wtNewas || ' '
 || wtDelta);

 opg_apis.find_cc_mapping_based('mypgGE$', wtClusters, wtUnDir,
 wtCluas, wtNewas, wtDelta, 8, 0, 'MYTBS', 'PDML=T');

 --
 -- logic to consume results in wtClusters
 -- e.g.:
 -- select /*+ parallel(8) */ count(distinct cluster_id)
 -- from mypg_clusters;

 -- cleanup all the working tables
 opg_apis.find_clusters_cleanup('mypgGE$', wtClusters, wtUnDir,
 wtCluas, wtNewas, wtDelta, '');

END;
/
```

## 9.27 OPG\_APIS.FIND\_CLUSTERS\_CLEANUP

**Format**

```
OPG_APIS.FIND_CLUSTERS_CLEANUP(
 edge_tab_name IN VARCHAR2,
 wt_clusters IN OUT VARCHAR2,
 wt_undir IN OUT VARCHAR2,
```

```

wt_cluas IN OUT VARCHAR2,
wt_newas IN OUT VARCHAR2,
wt_delta IN OUT VARCHAR2,
options IN VARCHAR2 DEFAULT NULL);

```

**Description**

Cleans up after running weakly connected components (WCC) cluster detection.

**Parameters****edge\_tab\_name**

Name of the property graph edge table.

**wt\_clusters**

A working table holding the final vertex cluster mappings. This table has two columns (VID NUMBER, CLUSTER\_ID NUMBER). Column VID stores the vertex ID values, and column CLUSTER\_ID stores the corresponding cluster ID values. Cluster ID values are long integers that can have gaps between them. If an empty name is specified, a new table will be generated, and its name will be returned.

**wt\_undir**

A working table holding an undirected version of the graph.

**wt\_cluas**

A working table holding current cluster assignments.

**wt\_newas**

A working table holding updated cluster assignments.

**wt\_delta**

A working table holding changes ("delta") in cluster assignments.

**options**

(Reserved for future use.)

**Usage Notes**

The property graph edge table must exist in the database.

**Examples**

The following example cleans up after performing doing cluster detection in a property graph named mypg.

```

EXECUTE OPG_APIS.FIND_CLUSTERS_CLEANUP('mypgGE$', wtClusters, wtUnDir, wtCluas,
wtNewas, wtDelta, null);

```

## 9.28 OPG\_APIS.FIND\_CLUSTERS\_PREP

**Format**

```

OPG_APIS.FIND_CLUSTERS_PREP (
 edge_tab_name IN VARCHAR2,
 wt_clusters IN OUT VARCHAR2,
 wt_undir IN OUT VARCHAR2,
 wt_cluas IN OUT VARCHAR2,

```



```

wt_newas IN OUT VARCHAR2,
wt_delta IN OUT VARCHAR2,
options IN VARCHAR2 DEFAULT NULL);

```

### Description

Prepares for running weakly connected components (WCC) cluster detection.

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table.

#### **wt\_clusters**

A working table holding the final vertex cluster mappings. This table has two columns (VID NUMBER, CLUSTER\_ID NUMBER). Column VID stores the vertex ID values, and column CLUSTER\_ID stores the corresponding cluster ID values. Cluster ID values are long integers that can have gaps between them.

If an empty name is specified, a new table will be generated, and its name will be returned.

#### **wt\_undir**

A working table holding an undirected version of the graph.

#### **wt\_cluas**

A working table holding current cluster assignments.

#### **wt\_newas**

A working table holding updated cluster assignments.

#### **wt\_delta**

A working table holding changes ("delta") in cluster assignments.

#### **options**

Additional settings for index creation.

### Usage Notes

The property graph edge table must exist in the database.

### Examples

The following example prepares for doing cluster detection in a property graph named `mypg`.

```

DECLARE
 wtClusters varchar2(200);
 wtUnDir varchar2(200);
 wtCluas varchar2(200);
 wtNewas varchar2(200);
 wtDelta varchar2(200);
BEGIN
 opg_apis.find_clusters_prep('mypgGE$', wtClusters, wtUnDir,
 wtCluas, wtNewas, wtDelta, '');
 dbms_output.put_line('working tables names ' || wtClusters || ' '
 || wtUnDir || ' ' || wtCluas || ' ' || wtNewas || ' '
 || wtDelta);
END;
/

```

## 9.29 OPG\_APIS.FIND\_SP

### Format

```
OPG_APIS.FIND_SP(
 edge_tab_name IN VARCHAR2,
 source IN NUMBER,
 dest IN NUMBER,
 exp_tab IN OUT VARCHAR2,
 dop IN INTEGER,
 stats_freq IN INTEGER DEFAULT 20000,
 path_output OUT VARCHAR2,
 weights_output OUT VARCHAR2,
 edge_tab_name IN VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL,
 scn IN NUMBER DEFAULT NULL);
```

### Description

Finds the shortest path between given source vertex and destination vertex in the property graph. It assumes each edge has a numeric weight property. (The actual edge property name is not significant.)

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table.

#### **source**

Source (start) vertex ID.

#### **dest**

Destination (end) vertex ID.

#### **exp\_tab**

Name of the expansion table to be used for shortest path calculations.

#### **dop**

Degree of parallelism for the operation.

#### **stats\_freq**

Frequency for collecting statistics on the table.

#### **path\_output**

The output shortest path. It consists of IDs of vertices on the shortest path, which are separated by the space character.

#### **weights\_output**

The output shortest path weights. It consists of weights of edges on the shortest path, which are separated by the space character.

#### **options**

Additional settings for the operation. An optional string with one or more (comma-separated) of the following values:

- CREATE\_UNDIRECTED=T
- REUSE\_UNDIRECTED\_TAB=T

**scn**

SCN for the edge table. It can be null.

**Usage Notes**

The property graph edge table must exist in the database, and the [OPG\\_APIS.FIND\\_SP\\_PREP](#) procedure must have already been called.

**Examples**

The following example prepares for shortest-path calculation, and then finds the shortest path from vertex 1 to vertex 35 in a property graph named `mypg`.

```
set serveroutput on
DECLARE
 w varchar2(2000);
 wtExp varchar2(2000);
 vPath varchar2(2000);
BEGIN
 opg_apis.find_sp_prep('mypgGE$', wtExp, null);
 opg_apis.find_sp('mypgGE$', 1, 35, wtExp, 1, 200000, vPath, w, null, null);
 dbms_output.put_line('Shortest path ' || vPath);
 dbms_output.put_line('Path weights ' || w);
END;
/
```

The output will be similar to the following. It shows one shortest path starting from vertex 1, to vertex 2, and finally to the destination vertex (35).

```
Shortest path 1 2 35
Path weights 3 2 1 1
```

## 9.30 OPG\_APIS.FIND\_SP\_CLEANUP

**Format**

```
OPG_APIS.FIND_SP_CLEANUP(
 edge_tab_name IN VARCHAR2,
 exp_tab IN OUT VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

**Description**

Cleans up after running one or more shortest path calculations.

**Parameters****edge\_tab\_name**

Name of the property graph edge table.

**exp\_tab**

Name of the expansion table used for shortest path calculations.

**options**

(Reserved for future use.)

**Usage Notes**

There is no need to call this procedure after each `OPG_APIS.FIND_SP` call. You can run multiple shortest path calculations before calling `OPG_APIS.FIND_SP_CLEANUP`.

**Examples**

The following example does cleanup work after doing shortest path calculations in a property graph named `mypg`.

```
EXECUTE OPG_APIS.FIND_SP_CLEANUP('mypgGE$', wtExpTab, null);
```

## 9.31 OPG\_APIS.FIND\_SP\_PREP

**Format**

```
OPG_APIS.FIND_SP_PREP(
 edge_tab_name IN VARCHAR2,
 exp_tab IN OUT VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

**Description**

Prepares for shortest path calculations.

**Parameters****edge\_tab\_name**

Name of the property graph edge table.

**exp\_tab**

Name of the expansion table to be used for shortest path calculations. If it is empty, an intermediate working table will be created and the table name will be returned in `exp_tab`.

**options**

Additional settings for the operation. An optional string with one or more (comma-separated) of the following values:

- `CREATE_UNDIRECTED=T`
- `REUSE_UNDIRECTED_TAB=T`

**Usage Notes**

The property graph edge table must exist in the database.

**Examples**

The following example does preparation work before doing shortest path calculations in a property graph named `mypg`

```
set serveroutput on
DECLARE
 wtExp varchar2(2000); -- name of working table for shortest path calculation
BEGIN
```

```
 opg_apis.find_sp_prep('mypgGE$', wtExp, null);
 dbms_output.put_line('Working table name ' || wtExp);
END;
/
```

The output will be similar to the following. (Your output may be different depending on the SQL session ID.)

```
Working table name "MYPGGE$$TWFS277"
```

## 9.32 OPG\_APIS.GET\_BUILD\_ID

### Format

```
OPG_APIS.GET_BUILD_ID() RETURN VARCHAR2;
```

### Description

Returns the current build ID of the Oracle Spatial and Graph property graph support, in YYYYMMDD format.

### Parameters

(None.)

### Usage Notes

(None.)

### Examples

The following example returns the current build ID of the Oracle Spatial and Graph property graph support.

```
SQL> SELECT OPG_APIS.GET_BUILD_ID() FROM DUAL;
```

```
OPG_APIS.GET_BUILD_ID()
```

```

20160606
```

## 9.33 OPG\_APIS.GET\_GEOMETRY\_FROM\_V\_COL

### Format

```
OPG_APIS.GET_GEOMETRY_FROM_V_COL(
 v IN NVARCHAR2,
 srid IN NUMBER DEFAULT 8307
) RETURN SDO_GEOMETRY;
```

### Description

Returns an SDO\_GEOMETRY object constructed using spatial data and optionally an SRID value.

## Parameters

### v

A String containing spatial data in serialized form.

### srid

SRID (coordinate system identifier) to be used in the resulting SDO\_GEOMETRY object. The default value is 8307, the Oracle Spatial SRID for the WGS 84 longitude/latitude coordinate system.

## Usage Notes

If there is incorrect syntax or a parsing error, this function returns NULL instead of generating an exception.

## Examples

The following examples show point, line, and polygon geometries.

```
SQL> select opg_apis.get_geometry_from_v_col('10.0 5.0',8307) from
dual;
```

```
OPG_APIS.GET_GEOMETRY_FROM_V_COL('10.05.0',8307) (SDO_GTYPE, SDO_SRID,
SDO_POINT(
```

```


SDO_GEOMETRY(2001, 8307, SDO_POINT_TYPE(10, 5, NULL), NULL, NULL)
```

```
SQL> select opg_apis.get_geometry_from_v_col('LINESTRING(30 10, 10
30, 40 40)',8307) from dual;
```

```
OPG_APIS.GET_GEOMETRY_FROM_V_COL('LINESTRING(3010,1030,4040)',8307)
(SDO_GTYPE, S
```

```


SDO_GEOMETRY(2002, 8307, NULL, SDO_ELEM_INFO_ARRAY(1, 2, 1),
SDO_ORDINATE_ARRAY(
30, 10, 10, 30, 40, 40))
```

```
SQL> select opg_apis.get_geometry_from_v_col('POLYGON((-83.6 34.1,
-83.6 34.3, -83.4 34.3, -83.4 34.1, -83.6 34.1))', 8307) from dual;
```

```
OPG_APIS.GET_GEOMETRY_FROM_V_COL('POLYGON((-83.634.1,-83.634.3,-83.434.
3,-83.434
```

```


SDO_GEOMETRY(2003, 8307, NULL, SDO_ELEM_INFO_ARRAY(1, 1003, 1),
SDO_ORDINATE_ARR
AY(-83.6, 34.1, -83.6, 34.3, -83.4, 34.3, -83.4, 34.1, -83.6, 34.1))
```

## 9.34 OPG\_APIS.GET\_GEOMETRY\_FROM\_V\_T\_COLS

### Format

```
OPG_APIS.GET_GEOMETRY_FROM_V_T_COLS(
 v IN NVARCHAR2,
 t IN INTEGER,
 srid IN NUMBER DEFAULT 8307
) RETURN SDO_GEOMETRY;
```

### Description

Returns an SDO\_GEOMETRY object constructed using spatial data, a type value, and optionally an SRID value.

### Parameters

#### v

A String containing spatial data in serialized form,

#### t

Value indicating the type of value represented by the v parameter. Must be 20. (A null value or any other value besides 20 returns a null SDO\_GEOMETRY object.)

#### srid

SRID (coordinate system identifier) to be used in the resulting SDO\_GEOMETRY object. The default value is 8307, the Oracle Spatial SRID for the WGS 84 longitude/latitude coordinate system.

### Usage Notes

If there is incorrect syntax or a parsing error, this function returns NULL instead of generating an exception.

### Examples

The following examples show point, line, and polygon geometries.

```
SQL> select opg_apis.get_geometry_from_v_t_cols('10.0 5.0', 20, 8307) from
dual;
```

```
OPG_APIS.GET_GEOMETRY_FROM_V_T_COLS('10.05.0',20,8307) (SDO_GTYPE, SDO_SRID,
SDO_

SDO_GEOMETRY(2001, 8307, SDO_POINT_TYPE(10, 5, NULL), NULL, NULL)
```

```
SQL> select opg_apis.get_geometry_from_v_t_cols('LINESTRING(30 10, 10 30, 40
40)', 20, 8307) from dual;
```

```
OPG_APIS.GET_GEOMETRY_FROM_V_T_COLS('LINESTRING(3010,1030,4040)',20,8307)
(SDO_GT


```

```
SDO_GEOMETRY(2002, 8307, NULL, SDO_ELEM_INFO_ARRAY(1, 2, 1),
SDO_ORDINATE_ARRAY(
30, 10, 10, 30, 40, 40))
```

```
SQL> select opg_apis.get_geometry_from_v_t_cols('POLYGON((-83.6
34.1, -83.6 34.3, -83.4 34.3, -83.4 34.1, -83.6 34.1))', 20, 8307)
from dual;
```

```
OPG_APIS.GET_GEOMETRY_FROM_V_T_COLS('POLYGON((-83.634.1,-83.634.3,-83.4
34.3,-83.
```

```

SDO_GEOMETRY(2003, 8307, NULL, SDO_ELEM_INFO_ARRAY(1, 1003, 1),
SDO_ORDINATE_ARR
AY(-83.6, 34.1, -83.6, 34.3, -83.4, 34.3, -83.4, 34.1, -83.6, 34.1))
```

## 9.35 OPG\_APIS.GET\_LATLONG\_FROM\_V\_COL

### Format

```
OPG_APIS.GET_LATLONG_FROM_V_COL(
 v IN NVARCHAR2,
 srid IN NUMBER DEFAULT 8307
) RETURN SDO_GEOMETRY;
```

### Description

Returns an SDO\_GEOMETRY object constructed using spatial data and optionally an SRID value.

### Parameters

#### v

A String containing spatial data in serialized form.

#### srid

SRID (coordinate system identifier) to be used in the resulting SDO\_GEOMETRY object. The default value is 8307, the Oracle Spatial SRID for the WGS 84 longitude/latitude coordinate system.

### Usage Notes

This function assumes that for each vertex in the geometry in the v parameter, the **first** number is the **latitude** value and the second number is the longitude value. (This is the reverse of the order in an SDO\_GEOMETRY object definition, where longitude is first and latitude is second).

If there is incorrect syntax or a parsing error, this function returns NULL instead of generating an exception.



## Examples

The following example returns a point SDO\_GEOMETRY object. Notice that the coordinate values of the input point are “swapped” in the returned SDO\_GEOMETRY object.

```
SQL> select opg_apis.get_latlong_from_v_col('5.1 10.0', 8307) from dual;

OPG_APIS.GET_LATLONG_FROM_V_COL('5.110.0',8307) (SDO_GTYPE, SDO_SRID,
SDO_POINT(X

SDO_GEOMETRY(2001, 8307, SDO_POINT_TYPE(10, 5.1, NULL), NULL, NULL)
```

## 9.36 OPG\_APIS.GET\_LATLONG\_FROM\_V\_T\_COLS

### Format

```
OPG_APIS.GET_LATLONG_FROM_V_T_COLS(
 v IN NVARCHAR2,
 t IN INTEGER,
 srid IN NUMBER DEFAULT 8307
) RETURN SDO_GEOMETRY;
```

### Description

Returns an SDO\_GEOMETRY object constructed using spatial data, a type value, and optionally an SRID value.

### Parameters

#### v

A String containing spatial data in serialized form.

#### t

Value indicating the type of value represented by the v parameter. Must be 20. (A null value or any other value besides 20 returns a null SDO\_GEOMETRY object.)

#### srid

SRID (coordinate system identifier) to be used in the resulting SDO\_GEOMETRY object. The default value is 8307, the Oracle Spatial SRID for the WGS 84 longitude/latitude coordinate system.

### Usage Notes

This function assumes that for each vertex in the geometry in the v parameter, the **first** number is the **latitude** value and the second number is the longitude value. (This is the reverse of the order in an SDO\_GEOMETRY object definition, where longitude is first and latitude is second).

If there is incorrect syntax or a parsing error, this function returns NULL instead of generating an exception.

## Examples

The following example returns a point SDO\_GEOMETRY object. Notice that the coordinate values of the input point are “swapped” in the returned SDO\_GEOMETRY object.

```
SQL> select opg_apis.get_latlong_from_v_t_cols('5.1 10.0',20,8307)
from dual;
```

```
OPG_APIS.GET_LATLONG_FROM_V_T_COLS('5.110.0',20,8307) (SDO_GTYPE,
SDO_SRID, SDO_P
```

```


SDO_GEOMETRY(2001, 8307, SDO_POINT_TYPE(10, 5.1, NULL), NULL, NULL)
```

## 9.37 OPG\_APIS.GET\_LONG\_LAT\_GEOMETRY

### Format

```
OPG_APIS.GET_LONG_LAT_GEOMETRY (
 x IN NUMBER,
 y IN NUMBER,
 srid IN NUMBER DEFAULT 8307
) RETURN SDO_GEOMETRY;
```

### Description

Returns an SDO\_GEOMETRY object constructed using X and Y point coordinate values, and optionally an SRID value.

### Parameters

#### x

The X (first coordinate) value in the SDO\_POINT\_TYPE element of the geometry definition.

#### y

The Y (second coordinate) value in the SDO\_POINT\_TYPE element of the geometry definition.

#### srid

SRID (coordinate system identifier) to be used in the resulting SDO\_GEOMETRY object. The default value is 8307, the Oracle Spatial SRID for the WGS 84 longitude/latitude coordinate system.

### Usage Notes

If there is incorrect syntax or a parsing error, this function returns NULL instead of generating an exception.

## Examples

The following example returns the geometry object for a point with X, Y coordinates 10.5, 5.0, and it uses 8307 as the SRID in the resulting geometry object.

```
SQL> select opg_apis.get_long_lat_geometry(10.0, 5.0, 8307) from dual;
```

```
OPG_APIS.GET_LONG_LAT_GEOMETRY(10.0,5.0,8307) (SDO_GTYPE, SDO_SRID,
SDO_POINT(X,
```

```


```

```
SDO_GEOMETRY(2001, 8307, SDO_POINT_TYPE(10, 5, NULL), NULL, NULL)
```

## 9.38 OPG\_APIS.GET\_LATLONG\_FROM\_V\_COL

### Format

```
OPG_APIS.GET_LATLONG_FROM_V_COL(
 v IN NVARCHAR2,
 srid IN NUMBER DEFAULT 8307
) RETURN SDO_GEOMETRY;
```

### Description

Returns an SDO\_GEOMETRY object constructed using spatial data and optionally an SRID value.

### Parameters

#### v

A String containing spatial data in serialized form.

#### srid

SRID (coordinate system identifier) to be used in the resulting SDO\_GEOMETRY object. The default value is 8307, the Oracle Spatial SRID for the WGS 84 longitude/latitude coordinate system.

### Usage Notes

This function assumes that for each vertex in the geometry in the *v* parameter, the **first** number is the **latitude** value and the second number is the longitude value. (This is the reverse of the order in an SDO\_GEOMETRY object definition, where longitude is first and latitude is second).

If there is incorrect syntax or a parsing error, this function returns NULL instead of generating an exception.

### Examples

The following example returns a point SDO\_GEOMETRY object. Notice that the coordinate values of the input point are “swapped” in the returned SDO\_GEOMETRY object.

```
SQL> select opg_apis.get_latlong_from_v_col('5.1 10.0', 8307) from dual;
```

```
OPG_APIS.GET_LATLONG_FROM_V_COL('5.110.0',8307) (SDO_GTYPE, SDO_SRID,
```

```
SDO_POINT(X

SDO_GEOMETRY(2001, 8307, SDO_POINT_TYPE(10, 5.1, NULL), NULL, NULL)
```

## 9.39 OPG\_APIS.GET\_LONGLAT\_FROM\_V\_T\_COLS

### Format

```
OPG_APIS.GET_LONGLAT_FROM_V_T_COLS(
 v IN NVARCHAR2,
 t IN INTEGER,
 srid IN NUMBER DEFAULT 8307
) RETURN SDO_GEOMETRY;
```

### Description

Returns an SDO\_GEOMETRY object constructed using spatial data, a type value, and optionally an SRID value.

### Parameters

#### v

A String containing spatial data in serialized form.

#### t

Value indicating the type of value represented by the v parameter. Must be 20. (A null value or any other value besides 20 returns a null SDO\_GEOMETRY object.)

#### srid

SRID (coordinate system identifier) to be used in the resulting SDO\_GEOMETRY object. The default value is 8307, the Oracle Spatial SRID for the WGS 84 longitude/latitude coordinate system.

### Usage Notes

If there is incorrect syntax or a parsing error, this function returns NULL instead of generating an exception.

### Examples

This function assumes that for each vertex in the geometry in the v parameter, the first number is the longitude value and the second number is the latitude value (which is the order in an SDO\_GEOMETRY object definition).

The following example returns a point SDO\_GEOMETRY object.

```
SQL> select opg_apis.get_longlat_from_v_t_cols('5.1 10.0',20,8307)
from dual;
```

```
OPG_APIS.GET_LATLONG_FROM_V_T_COLS('5.110.0',20,8307) (SDO_GTYPE,
SDO_SRID, SDO_P
```

```


```

```
SDO_GEOMETRY(2001, 8307, SDO_POINT_TYPE(5.1, 10, NULL), NULL, NULL)
```

## 9.40 OPG\_APIS.GET\_SCN

### Format

```
OPG_APIS.GET_SCN() RETURN NUMBER;
```

### Description

Returns the SCN (system change number) of the Oracle Spatial and Graph property graph support, in YYYYMMDD format.

#### Note:

Effective with Release 20.3, the OPG\_APIS.GET\_SCN function is **deprecated**. Instead, to retrieve the current SCN (system change number), use the DBMS\_FLASHBACK.GET\_SYSTEM\_CHANGE\_NUMBER function:

```
SELECT dbms_flashback.get_system_change_number FROM DUAL;
```

### Parameters

(None.)

### Usage Notes

The SCN value is incremented after each commit.

### Examples

The following example returns the current build ID of the Oracle Spatial and Graph property graph support.

```
SQL> SELECT OPG_APIS.GET_SCN() FROM DUAL;
```

```
OPG_APIS.GET_SCN()

1478701
```

## 9.41 OPG\_APIS.GET\_VERSION

### Format

```
OPG_APIS.GET_VERSION() RETURN VARCHAR2;
```

### Description

Returns the current version of the Oracle Spatial and Graph property graph support.

**Parameters**

(None.)

**Usage Notes**

(None.)

**Examples**

The following example returns the current version of the Oracle Spatial and Graph property graph support.

```
SQL> SELECT OPG_APIS.GET_VERSION() FROM DUAL;
```

```
OPG_APIS.GET_VERSION()
```

```

12.2.0.1 P1
```

## 9.42 OPG\_APIS.GET\_WKTGEOMETRY\_FROM\_V\_COL

**Format**

```
OPG_APIS.GET_WKTGEOMETRY_FROM_V_COL(
 v IN NVARCHAR2,
 srid IN NUMBER DEFAULT NULL
) RETURN SDO_GEOMETRY;
```

**Description**

Returns an SDO\_GEOMETRY object based on a geometry in WKT (well known text) form and optionally an SRID.

**Parameters****v**

A String containing spatial data in serialized form.

**srid**

SRID (coordinate system identifier) to be used in the resulting SDO\_GEOMETRY object. The default value is 8307, the Oracle Spatial SRID for the WGS 84 longitude/latitude coordinate system.

**Usage Notes**

If there is incorrect syntax or a parsing error, this function returns NULL instead of generating an exception.

**Examples**

The following statements return a point geometry and a line string geometry

```
SQL> select opg_apis.get_wktgeometry_from_v_col('POINT(10.0 5.1)',
8307) from dual;
```

```
OPG_APIS.GET_WKTGEOMETRY_FROM_V_COL('POINT(10.05.1)',8307) (SDO_GTYPE,
SDO_SRID,
```

```


SDO_GEOMETRY(2001, 8307, SDO_POINT_TYPE(10, 5.1, NULL), NULL, NULL)

SQL> select opg_apis.get_wktgeometry_from_v_col('LINESTRING(30 10, 10 30, 40
40)',8307) from dual;

OPG_APIS.GET_WKTGEOMETRY_FROM_V_COL('LINESTRING(3010,1030,4040)',8307)
(SDO_GTYPE

SDO_GEOMETRY(2002, 8307, NULL, SDO_ELEM_INFO_ARRAY(1, 2, 1),
SDO_ORDINATE_ARRAY(
30, 10, 10, 30, 40, 40))

```

## 9.43 OPG\_APIS.GET\_WKTGEOMETRY\_FROM\_V\_T\_COLS

### Format

```

OPG_APIS.GET_WKTGEOMETRY_FROM_V_T_COLS(
 v IN NVARCHAR2,
 t IN INTEGER,
 srid IN NUMBER DEFAULT NULL
) RETURN SDO_GEOMETRY;

```

### Description

Returns an SDO\_GEOMETRY object based on a geometry in WKT (well known text) form, a type value, and optionally an SRID.

### Parameters

#### v

A String containing spatial data in serialized form.

#### t

Value indicating the type of value represented by the v parameter. Must be 20. (A null value or any other value besides 20 returns a null SDO\_GEOMETRY object.)

#### srid

SRID (coordinate system identifier) to be used in the resulting SDO\_GEOMETRY object. The default value is 8307, the Oracle Spatial SRID for the WGS 84 longitude/latitude coordinate system.

### Usage Notes

If there is incorrect syntax or a parsing error, this function returns NULL instead of generating an exception.

### Examples

The following statements return a point geometry and a polygon geometry

```

SQL> select opg_apis.get_wktgeometry_from_v_t_cols('POINT(10.0
5.1)',20,8307) from dual;

```

```

OPG_APIS.GET_WKTGEOMETRY_FROM_V_T_COLS('POINT(10.05.1)',20,8307)
(SDO_GTYPE, SDO_

SDO_GEOMETRY(2001, 8307, SDO_POINT_TYPE(10, 5.1, NULL), NULL, NULL)

SQL> select opg_apis.get_wktgeometry_from_v_t_cols('POLYGON((-83.6
34.1, -83.6 34.3, -83.4 34.3, -83.4 34.1, -83.6 34.1))',20,8307) from
dual;

OPG_APIS.GET_WKTGEOMETRY_FROM_V_T_COLS('POLYGON((-83.634.1,-83.634.3,-8
3.434.3,-

SDO_GEOMETRY(2003, 8307, NULL, SDO_ELEM_INFO_ARRAY(1, 1003, 1),
SDO_ORDINATE_ARR
AY(-83.6, 34.1, -83.6, 34.3, -83.4, 34.3, -83.4, 34.1, -83.6, 34.1))

```

## 9.44 OPG\_APIS.GRANT\_ACCESS

### Format

```

OPG_APIS.GRANT_ACCESS(
 graph_owner IN VARCHAR2,
 graph_name IN VARCHAR2,
 other_user IN VARCHAR2,
 privilege IN VARCHAR2);

```

### Description

Grants access privileges on a property graph to another database user.

### Parameters

#### **graph\_owner**

Owner of the property graph.

#### **graph\_name**

Name of the property graph.

#### **other\_user**

Name of the database user to which one or more access privileges will be granted.

#### **privilege**

A string of characters indicating privileges: R for read, S for select, U for update, D for delete, I for insert, A for all. Do not use commas or any other delimiter.

If you specify A, do not specify any other values because A includes all access privileges.

### Usage Notes

(None.)



## Examples

The following example grants read and select (RS) privileges on the `mypg` property graph owned by database user SCOTT to database user PGUSR. It then connects as PGUSR and queries the `mypg` vertex table in the SCOTT schema.

```
CONNECT scott/<password>

EXECUTE OPG_APIS.GRANT_ACCESS('scott', 'mypg', 'pgusr', 'RS');

CONNECT pgusr/<password>

SELECT count(1) from scott.mypgVT$;
```

17

## 9.45 OPG\_APIS.IMP\_EDGE\_TAB\_STATS

### Format

```
OPG_APIS.IMP_EDGE_TAB_STATS (
 graph_name IN VARCHAR2,
 stattab IN VARCHAR2,
 statid IN VARCHAR2 DEFAULT NULL,
 cascade IN BOOLEAN DEFAULT TRUE,
 statown IN VARCHAR2 DEFAULT NULL,
 no_invalidate BOOLEAN DEFAULT FALSE,
 force BOOLEAN DEFAULT FALSE,
 stat_category IN VARCHAR2 DEFAULT 'OBJECT_STATS');
```

### Description

Retrieves statistics for the given property graph edge table (GE\$) from the user statistics table identified by `stattab` and stores them in the dictionary. If `cascade` is `TRUE`, all index statistics associated with the specified table are also imported.

### Parameters

**graph\_name**

Name of the property graph.

**stattab**

Name of the statistics table.

**statid**

Optional identifier to associate with these statistics within `stattab`.

**cascade**

If `TRUE`, column and index statistics are exported.

**statown**

Schema containing `stattab`.

**no\_invalidate**

If `TRUE`, does not invalidate the dependent cursors. If `FALSE`, invalidates the dependent cursors immediately. If `DBMS_STATS.AUTO_INVALIDATE` (the usual default) is in effect, Oracle Database decides when to invalidate dependent cursors.

**force**

If `TRUE`, performs the operation even if the statistics are locked.

**stat\_category**

Specifies what statistics to export, using a comma to separate values. The supported values are `'OBJECT_STATS'` (the default: table statistics, column statistics, and index statistics) and `'SYNOPSIS'` (auxiliary statistics created when statistics are incrementally maintained).

**Usage Notes**

(None.)

**Examples**

The following example creates a statistics table, exports into this table the edge table statistics, issues a query to count the relevant rows for the newly created statistics, and finally imports the statistics back.

```
EXECUTE OPG_APIS.CREATE_STAT_TABLE('mystat',null);

EXECUTE OPG_APIS.EXP_EDGE_TAB_STATS('mypg', 'mystat', 'edge_stats_id_1', true,
null, 'OBJECT_STATS');

SELECT count(1) FROM mystat WHERE statid='EDGE_STATS_ID_1';

 153

EXECUTE OPG_APIS.IMP_EDGE_TAB_STATS('mypg', 'mystat', 'edge_stats_id_1', true,
null, false, true, 'OBJECT_STATS');
```

## 9.46 OPG\_APIS.IMP\_VERTEX\_TAB\_STATS

**Format**

```
OPG_APIS.IMP_VERTEX_TAB_STATS(
 graph_name IN VARCHAR2,
 stattab IN VARCHAR2,
 statid IN VARCHAR2 DEFAULT NULL,
 cascade IN BOOLEAN DEFAULT TRUE,
 statown IN VARCHAR2 DEFAULT NULL,
 no_invalidate BOOLEAN DEFAULT FALSE,
 force BOOLEAN DEFAULT FALSE,
 stat_category IN VARCHAR2 DEFAULT 'OBJECT_STATS');
```

**Description**

Retrieves statistics for the given property graph vertex table (VT\$) from the user statistics table identified by `stattab` and stores them in the dictionary. If `cascade` is `TRUE`, all index statistics associated with the specified table are also imported.

## Parameters

### **graph\_name**

Name of the property graph.

### **stattab**

Name of the statistics table.

### **statid**

Optional identifier to associate with these statistics within `stattab`.

### **cascade**

If `TRUE`, column and index statistics are exported.

### **statown**

Schema containing `stattab`.

### **no\_invalidate**

If `TRUE`, does not invalidate the dependent cursors. If `FALSE`, invalidates the dependent cursors immediately. If `DBMS_STATS.AUTO_INVALIDATE` (the usual default) is in effect, Oracle Database decides when to invalidate dependent cursors.

### **force**

If `TRUE`, performs the operation even if the statistics are locked.

### **stat\_category**

Specifies what statistics to export, using a comma to separate values. The supported values are `'OBJECT_STATS'` (the default: table statistics, column statistics, and index statistics) and `'SYNOPSIS'` (auxiliary statistics created when statistics are incrementally maintained).

## Usage Notes

(None.)

## Examples

The following example creates a statistics table, exports into this table the vertex table statistics, issues a query to count the relevant rows for the newly created statistics, and finally imports the statistics back.

```
EXECUTE OPG_APIS.CREATE_STAT_TABLE('mystat',null);

EXECUTE OPG_APIS.EXP_VERTEX_TAB_STATS('mypg', 'mystat', 'vertex_stats_id_1', true,
null, 'OBJECT_STATS');

SELECT count(1) FROM mystat WHERE statid='VERTEX_STATS_ID_1';

108

EXECUTE OPG_APIS.IMP_VERTEX_TAB_STATS('mypg', 'mystat', 'vertex_stats_id_1', true,
null, false, true, 'OBJECT_STATS');
```

## 9.47 OPG\_APIS.PR

### Format

```
OPG_APIS.PR(
 edge_tab_name IN VARCHAR2,
 d IN NUMBER DEFAULT 0.85,
 num_iterations IN NUMBER DEFAULT 10,
 convergence IN NUMBER DEFAULT 0.1,
 dop IN INTEGER DEFAULT 4,
 wt_node_pr IN OUT VARCHAR2,
 wt_node_nextpr IN OUT VARCHAR2,
 wt_edge_tab_deg IN OUT VARCHAR2,
 wt_delta IN OUT VARCHAR2,
 tablespace IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL,
 num_vertices OUT NUMBER);
```

### Description

Prepares for page rank calculations.

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table.

#### **d**

Damping factor.

#### **num\_iterations**

Number of iterations for calculating the page rank values.

#### **convergence**

A threshold. If the difference between the page rank value of the current iteration and next iteration is lower than this threshold, then computation stops.

#### **dop**

Degree of parallelism for the operation.

#### **wt\_node\_pr**

Name of the working table to hold the page rank values of the vertices.

#### **wt\_node\_pr**

Name of the working table to hold the page rank values of the vertices.

#### **wt\_node\_next\_pr**

Name of the working table to hold the page rank values of the vertices in the next iteration.

#### **wt\_edge\_tab\_deg**

Name of the working table to hold edges and node degree information.

#### **wt\_delta**

Name of the working table to hold information about some special vertices.

**tablespace**

Name of the tablespace to hold all the graph data and index data.

**options**

Additional settings for the operation. An optional string with one or more (comma-separated) of the following values:

- CREATE\_UNDIRECTED=T
- REUSE\_UNDIRECTED\_TAB=T

**num\_vertices**

Number of vertices processed by the page rank calculation.

**Usage Notes**

The property graph edge table must exist in the database, and the [OPG\\_APIS.PR\\_PREP](#) procedure must have been called.

**Examples**

The following example performs preparation, and then calculates the page rank value of vertices in a property graph named `mypg`.

```
set serveroutput on
DECLARE
 wt_pr varchar2(2000); -- name of the table to hold PR value of the current
iteration
 wt_npr varchar2(2000); -- name of the table to hold PR value for the next iteration
 wt3 varchar2(2000);
 wt4 varchar2(2000);
 wt5 varchar2(2000);
 n_vertices number;
BEGIN
 wt_pr := 'mypgPR';
 opg_apis.pr_prep('mypgGE$', wt_pr, wt_npr, wt3, wt4, null);
 dbms_output.put_line('Working table names ' || wt_pr
 || ', wt_npr ' || wt_npr || ', wt3 ' || wt3 || ', wt4 ' || wt4);
 opg_apis.pr('mypgGE$', 0.85, 10, 0.01, 4, wt_pr, wt_npr, wt3, wt4, 'SYSAUX', null,
n_vertices)
;
END;
/
```

The output will be similar to the following.

```
Working table names "MYPGPR", wt_npr "MYPGGE$$TWPRX277", wt3
"MYPGGE$$TWPRE277", wt4 "MYPGGE$$TWPRD277"
```

The calculated page rank value is stored in the `mypgpr` table which has the following definition and data.

```
SQL> desc mypgpr;
Name Null? Type

NODE NOT NULL NUMBER
PR NUMBER
C NUMBER

SQL> select node, pr from mypgpr;
```

NODE	PR
101	.1925
201	.2775
102	.1925
104	.74383125
105	.313625
103	.1925
100	.15
200	.15

## 9.48 OPG\_APIS.PR\_CLEANUP

### Format

```
OPG_APIS.PR_CLEANUP(
 edge_tab_name IN VARCHAR2,
 wt_node_pr IN OUT VARCHAR2,
 wt_node_nextpr IN OUT VARCHAR2,
 wt_edge_tab_deg IN OUT VARCHAR2,
 wt_delta IN OUT VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Performs cleanup after performing page rank calculations.

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table.

#### **wt\_node\_pr**

Name of the working table to hold the page rank values of the vertices.

#### **wt\_node\_next\_pr**

Name of the working table to hold the page rank values of the vertices in the next iteration.

#### **wt\_edge\_tab\_deg**

Name of the working table to hold edges and node degree information.

#### **wt\_delta**

Name of the working table to hold information about some special vertices.

#### **options**

Additional settings for the operation. An optional string with one or more (comma-separated) of the following values:

- CREATE\_UNDIRECTED=T
- REUSE\_UNDIRECTED\_TAB=T

### Usage Notes

You do not need to do cleanup after each call to the [OPG\\_APIS.PR](#) procedure. You can run several page rank calculations before calling the [OPG\\_APIS.PR\\_CLEANUP](#) procedure.

## Examples

The following example does the cleanup work after running page rank calculations in a property graph named `mypg`.

```
EXECUTE OPG_APIS.PR_CLEANUP('mypgGE$', wt_pr, wt_npr, wt3, wt4, null);
```

## 9.49 OPG\_APIS.PR\_PREP

### Format

```
OPG_APIS.PR_PREP(
 edge_tab_name IN VARCHAR2,
 wt_node_pr IN OUT VARCHAR2,
 wt_node_nextpr IN OUT VARCHAR2,
 wt_edge_tab_deg IN OUT VARCHAR2,
 wt_delta IN OUT VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Prepares for page rank calculations.

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table.

#### **wt\_node\_pr**

Name of the working table to hold the page rank values of the vertices.

#### **wt\_node\_next\_pr**

Name of the working table to hold the page rank values of the vertices in the next iteration.

#### **wt\_edge\_tab\_deg**

Name of the working table to hold edges and node degree information.

#### **wt\_delta**

Name of the working table to hold information about some special vertices.

#### **options**

Additional settings for the operation. An optional string with one or more (comma-separated) of the following values:

- `CREATE_UNDIRECTED=T`
- `REUSE_UNDIRECTED_TAB=T`

### Usage Notes

The property graph edge table must exist in the database.

### Examples

The following example does the preparation work before running page rank calculations in a property graph named `mypg`.

```

set serveroutput on
DECLARE
 wt_pr varchar2(2000); -- name of the table to hold PR value of the current
iteration
 wt_npr varchar2(2000); -- name of the table to hold PR value for the next
iteration
 wt3 varchar2(2000);
 wt4 varchar2(2000);
 wt5 varchar2(2000);
BEGIN
 wt_pr := 'mypgPR';
 opg_apis.pr_prep('mypgGE$', wt_pr, wt_npr, wt3, wt4, null);
 dbms_output.put_line('Working table names ' || wt_pr
 || ', wt_npr ' || wt_npr || ', wt3 ' || wt3 || ', wt4 ' || wt4);
END;
/

```

The output will be similar to the following.

```

Working table names "MYPGPR", wt_npr "MYPGGE$$TWPRX277", wt3
"MYPGGE$$TWPRE277", wt4 "MYPGGE$$TWPRD277"

```

## 9.50 OPG\_APIS.PREPARE\_TEXT\_INDEX

### Format

```
OPG_APIS.PREPARE_TEXT_INDEX();
```

### Description

Performs preparatory work needed before a text index can be created on any NVARCHAR2 columns.

### Parameters

None.

### Usage Notes

You must have the ALTER SESSION to run this procedure.

### Examples

The following example performs preparatory work needed before a text index can be created on any NVARCHAR2 columns.

```
EXECUTE OPG_APIS.PREPARE_TEXT_INDEX();
```

## 9.51 OPG\_APIS.RENAME\_PG

### Format

```
OPG_APIS.RENAME_PG(
 graph_name IN VARCHAR2,
 new_graph_name IN VARCHAR2);
```



**Description**

Renames a property graph.

**Parameters****graph\_name**

Name of the property graph.

**new\_graph\_name**

New name for the property graph.

**Usage Notes**

The `graph_name` property graph must exist in the database.

**Examples**

The following example changes the name of a property graph named `mypg` to `mynewpg`.

```
EXECUTE OPG_APIS.RENAME_PG('mypg', 'mynewpg');
```

## 9.52 OPG\_APIS.SPARSIFY\_GRAPH

**Format**

```
OPG_APIS.SPARSIFY_GRAPH(
 edge_tab_name IN VARCHAR2,
 threshold IN NUMBER DEFAULT 0.5,
 min_keep IN INTEGER DEFAULT 1,
 dop IN INTEGER DEFAULT 4,
 wt_out_tab IN OUT VARCHAR2,
 wt_und_tab IN OUT VARCHAR2,
 wt_hsh_tab IN OUT VARCHAR2,
 wt_mch_tab IN OUT VARCHAR2,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL);
```

**Description**

Performs sparsification (edge trimming) for a property graph edge table.

**Parameters****edge\_tab\_name**

Name of the property graph edge table (GE\$).

**threshold**

A numeric value controlling how much sparsification needs to be performed. The lower the value, the more edges will be removed. Some typical values are: 0.1, 0.2, ..., 0.5

**min\_keep**

A positive integer indicating at least how many adjacent edges should be kept for each vertex. A recommended value is 1.

**dop**

Degree of parallelism for the operation.

**wt\_out\_tab**

A working table to hold the output, a sparsified graph.

**wt\_und\_tab**

A working table to hold the undirected version of the original graph.

**wt\_hsh\_tab**

A working table to hold the min hash values of the graph.

**wt\_mch\_tab**

A working table to hold matching count of min hash values.

**tbs**

A working table to hold the working table data.

**options**

Additional settings for operation. An optional string with one or more (comma-separated) of the following values:

- 'INMEMORY=T' is an option for creating the schema tables with an 'inmemory' clause.
- 'IMC\_MC\_B=T' creates the schema tables with an INMEMORY MEMCOMPRESS BASIC clause.

**Usage Notes**

The CREATE TABLE privilege is required to call this procedure.

The sparsification algorithm used is a min hash based local sparsification. See "Local graph sparsification for scalable clustering", Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data: <https://cs.uwaterloo.ca/~tozsu/courses/CS848/W15/presentations/ElbagouryPresentation-2.pdf>

Sparsification only involves the topology of a graph. None of the properties (K/V) are relevant.

**Examples**

The following example does the preparation work for the edges table of `mypg`, prints out the working table names, and runs sparsification. The output, a sparsified graph, is stored in a table named `LEAN_PG`, which has two columns, `SVID` and `DVID`.

```
SQL> set serveroutput on
DECLARE
 my_lean_pg varchar2(100) := 'lean_pg'; -- output table
 wt2 varchar2(100);
 wt3 varchar2(100);
 wt4 varchar2(100);
BEGIN
 opg_apis.sparsify_graph_prep('mypgGE$', my_lean_pg, wt2, wt3, wt4, null);
 dbms_output.put_line('wt2 ' || wt2 || ', wt3 ' || wt3 || ', wt4 ' || wt4);

 opg_apis.sparsify_graph('mypgGE$', 0.5, 1, 4, my_lean_pg, wt2, wt3, wt4,
'SEMTS', null);
END;
/
```

```
wt2 "MYPGGE$$TWSPA275", wt3 "MYPGGE$$TWSPA275", wt4 "MYPGGE$$TWSPAM275"
```

```
SQL> describe lean_pg;
```

Name	Null?	Type
SVID		NUMBER
DVID		NUMBER

## 9.53 OPG\_APIS.SPARSIFY\_GRAPH\_CLEANUP

### Format

```
OPG_APIS.SPARSIFY_GRAPH_CLEANUP(
 edge_tab_name IN VARCHAR2,
 wt_out_tab IN OUT VARCHAR2,
 wt_und_tab IN OUT VARCHAR2,
 wt_hsh_tab IN OUT VARCHAR2,
 wt_mch_tab IN OUT VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Cleans up after sparsification (edge trimming) for a property graph edge table.

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table (GE\$).

#### **wt\_out\_tab**

A working table to hold the output, a sparsified graph.

#### **wt\_und\_tab**

A working table to hold the undirected version of the original graph.

#### **wt\_hsh\_tab**

A working table to hold the min hash values of the graph.

#### **wt\_mch\_tab**

A working table to hold matching count of min hash values.

#### **tbs**

A working table to hold the working table data

#### **options**

(Reserved for future use.)

### Usage Notes

The working tables will be dropped after the operation completes.

### Examples

The following example does the preparation work for the edges table of `mypg`, prints out the working table names, runs sparsification, and then performs cleanup.

```
SQL> set serveroutput on
DECLARE
 my_lean_pg varchar2(100) := 'lean_pg';
 wt2 varchar2(100);
 wt3 varchar2(100);
 wt4 varchar2(100);
BEGIN
 opg_apis.sparsify_graph_prep('mypgGE$', my_lean_pg, wt2, wt3, wt4, null);
 dbms_output.put_line('wt2 ' || wt2 || ', wt3 ' || wt3 || ', wt4 ' || wt4);

 opg_apis.sparsify_graph('mypgGE$', 0.5, 1, 4, my_lean_pg, wt2, wt3, wt4,
'SEMTS', null);

 -- Add logic here to consume SVID, DVID in LEAN_PG table
 --

 -- cleanup
 opg_apis.sparsify_graph_cleanup('mypgGE$', my_lean_pg, wt2, wt3, wt4, null);
END;
/
```

## 9.54 OPG\_APIS.SPARSIFY\_GRAPH\_PREP

### Format

```
OPG_APIS.SPARSIFY_GRAPH_PREP(
 edge_tab_name IN VARCHAR2,
 wt_out_tab IN OUT VARCHAR2,
 wt_und_tab IN OUT VARCHAR2,
 wt_hsh_tab IN OUT VARCHAR2,
 wt_mch_tab IN OUT VARCHAR2,
 options IN VARCHAR2 DEFAULT NULL);
```

### Description

Prepares working table names that are necessary to run sparsification for a property graph edge table.

### Parameters

#### **edge\_tab\_name**

Name of the property graph edge table (GE\$).

#### **wt\_out\_tab**

A working table to hold the output, a sparsified graph.

#### **wt\_und\_tab**

A working table to hold the undirected version of the original graph.

#### **wt\_hsh\_tab**

A working table to hold the min hash values of the graph.

#### **wt\_mch\_tab**

A working table to hold the matching count of min hash values.

**options**

Additional settings for operation. An optional string with one or more (comma-separated) of the following values:

- 'INMEMORY=T' is an option for creating the schema tables with an 'inmemory' clause.
- 'IMC\_MC\_B=T' creates the schema tables with an INMEMORY MEMCOMPRESS BASIC clause.

**Usage Notes**

The sparsification algorithm used is a min hash based local sparsification. See "Local graph sparsification for scalable clustering", Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data: <https://cs.uwaterloo.ca/~tozsu/courses/CS848/W15/presentations/ElbagouryPresentation-2.pdf>

**Examples**

The following example does the preparation work for the edges table of `mypg` and prints out the working table names.

```
set serveroutput on

DECLARE
 my_lean_pg varchar2(100) := 'lean_pg';
 wt2 varchar2(100);
 wt3 varchar2(100);
 wt4 varchar2(100);
BEGIN
 opg_apis.sparsify_graph_prep('mypgGE$', my_lean_pg, wt2, wt3, wt4, null);
 dbms_output.put_line('wt2 ' || wt2 || ', wt3 ' || wt3 || ', wt4 ' || wt4);
END;
/
```

The output may be similar to the following.

```
wt2 "MYPGGE$$TWSPA275", wt3 "MYPGGE$$TWSPA275", wt4 "MYPGGE$$TWSPAM275"
```

# 10

## OPG\_GRAPHOP Package Subprograms

The OPG\_GRAPHOP package contains subprograms for various operations on property graphs in an Oracle database.

To use the subprograms in this chapter, you must understand the conceptual and usage information in earlier chapters of this book.

This chapter provides reference information about the subprograms, in alphabetical order.

- [OPG\\_GRAPHOP.POPULATE\\_SKELETON\\_TAB](#)

### 10.1 OPG\_GRAPHOP.POPULATE\_SKELETON\_TAB

#### Format

```
OPG_GRAPHOP.POPULATE_SKELETON_TAB(
 graph IN VARCHAR2,
 dop IN INTEGER DEFAULT 4,
 tbs IN VARCHAR2 DEFAULT NULL,
 options IN VARCHAR2 DEFAULT NULL);
```

#### Description

Populates the skeleton table (<graph-name>GT\$). By default, any existing content in the skeleton table is truncated (removed) before the table is populated.

#### Parameters

##### **graph**

Name of the property graph.

##### **dop**

Degree of parallelism for the operation.

##### **tbs**

Name of the tablespace to hold the index data for the skeleton table.

##### **options**

Options that can be used to customize the populating of the skeleton table. (One or more, comma separated.)

- 'KEEP\_DATA=T' causes any existing table not to be removed before the table is populated. New rows are added after the existing ones.
- 'PDML=T' skips the default index creation.

#### Usage Notes

You must have the CREATE TABLE and CREATE INDEX privileges to call this procedure.

There is a unique index constraint on EID column of the skeleton table (GE\$). So if you specify the `KEEP_DATA=T` option and if the new data overlaps with existing one, then the unique key constraint will be violated, resulting in an error.

### Examples

The following example populates the skeleton table of the property graph named `mypg`.

```
EXECUTE OPG_GRAPHOP.POPULATE_SKELETON_TAB('mypg',4, 'pgts', 'PDML=T');
```

# Part II

## In-Memory Graph Server (PGX) Advanced User Guide

Part II provides in-depth information on using the in-memory graph server (PGX) for advanced users.

Part II contains the following chapters:

- [Configuring the In-Memory Graph Server \(PGX\)](#)  
This chapter explains the configuration options for the in-memory graph server (PGX) and the graph client.
- [Graphs Management](#)  
You can load, publish, store and delete graphs
- [Namespaces and Sharing](#)  
The in-memory graph server (PGX) supports separate namespaces that help you to organize your entities.
- [PGX Programming Guides](#)  
You can avail all the PGX functionalities through asynchronous Java APIs. Each asynchronous method has a synchronous equivalent, which blocks the caller thread until the server produces a response.
- [Working with Files Using the Graph Server \(PGX\)](#)  
This chapter describes in detail about working with different file formats to perform various actions like loading, storing or exporting a graph using the Graph Server (PGX).
- [Log Management in the Graph Server \(PGX\)](#)  
The graph server (PGX) internally uses the SLF4J interface with Log4j as the default logger implementation.



# 11

## Configuring the In-Memory Graph Server (PGX)

This chapter explains the configuration options for the in-memory graph server (PGX) and the graph client.

These options can be configured in `/etc/oracle/graph/pgx.conf` where the graph server is installed, or passed to the graph server programmatically.

- [Configuration Parameters for the Graph Server \(PGX\) Engine](#)  
You can configure the graph server (PGX) engine and the PGX run-time library by assigning a single JSON file to the in-memory graph server (PGX) at startup.
- [Configuration Parameters for Connecting to the Graph Server \(PGX\)](#)  
You can configure the graph server (PGX) to use the required options at startup.
- [Configuration Parameters for the Graph Client](#)  
You can configure the PGX graph client. All the parameters are available as command-line options also.

### 11.1 Configuration Parameters for the Graph Server (PGX) Engine

You can configure the graph server (PGX) engine and the PGX run-time library by assigning a single JSON file to the in-memory graph server (PGX) at startup.

This file includes the parameters shown in the following table.

To specify the configuration file, see [Specifying the Configuration File to the In-Memory Graph Server \(PGX\)](#).

#### Note:

- Relative paths in parameter values are always resolved relative to the parent directory of the configuration file in which they are specified. For example, if the configuration file is `/pgx/conf/pgx.conf`, then the file path `graph-configs/my-graph.bin.json` inside that file would be resolved to `/pgx/conf/graph-configs/my-graph.bin.json`.
- The parameter default values are optimized to deliver the best performance across a wide set of algorithms. Depending on your workload, you may be able to improve performance further by experimenting with different strategies, sizes, and thresholds.

**Table 11-1 Configuration Parameters for the Graph Server (PGX) Engine**


Parameter	Type	Description	Default
admin_request_cache_timeout	integer	After how many seconds admin request results get removed from the cache. Requests which are not done or not yet consumed are excluded from this timeout. Note: This is only relevant if PGX is deployed as a webapp.	60
allow_idle_timeout_overwrite	boolean	If true, sessions can overwrite the default idle timeout.	true
allow_override_scheduling_information	boolean	If true, allow all users to override scheduling information like task weight, task priority, and number of threads	true
allow_task_timeout_overwrite	boolean	If true, sessions can overwrite the default task timeout.	true
allow_user_auto_refresh	boolean	If true, users may enable auto refresh for graphs they load. If false, only graphs mentioned in <code>preload_graphs</code> can have auto refresh enabled.	false
allowed_remote_loading_locations	array of strings	Allow loading graphs into the PGX engine from remote locations (http, https, ftp, ftps, s3, hdfs). If empty, as by default, no remote location is allowed. If "*" is specified in the array, all remote locations are allowed. Only the value "*" is currently supported. Note that pre-loaded graphs are loaded from any location, regardless of the value of this setting.	[]
<div style="border: 1px solid #ccc; background-color: #fff9c4; padding: 10px; display: inline-block;">  <b>WARNING:</b>            This parameter reduces security and therefore use it only when needed.         </div>			
basic_scheduler_config	object	Configuration parameters for the fork join pool backend.	null
bfs_iterate_que_task_size	integer	Task size for BFS iterate QUE phase.	128
bfs_threshold_parent_read_based	number	Threshold of BFS traversal level items to switch to parent-read-based visiting strategy.	0.05
bfs_threshold_read_based	integer	Threshold of BFS traversal level items to switch to read-based visiting strategy.	1024
bfs_threshold_single_threaded	integer	Until what number of BFS traversal level items vertices are visited single-threaded.	128
character_set	string	Standard character set to use throughout PGX. UTF-8 is the default. Note: Some formats may not be compatible.	utf-8

Table 11-1 (Cont.) Configuration Parameters for the Graph Server (PGX) Engine

Parameter	Type	Description	Default
<code>cni_diff_factor_default</code>	integer	Default diff factor value used in the common neighbor iterator implementations.	8
<code>cni_small_default</code>	integer	Default value used in the common neighbor iterator implementations, to indicate below which threshold a subarray is considered small.	128
<code>cni_stop_recursion_default</code>	integer	Default value used in the common neighbor iterator implementations, to indicate the minimum size where the binary search approach is applied.	96
<code>dfs_threshold_large</code>	integer	Value that determines at which number of visited vertices the DFS implementation will switch to data structures that are optimized for larger numbers of vertices.	4096
<code>enable_csrf_token_checks</code>	boolean	If true, the PGX webapp will verify the Cross-Site Request Forgery (CSRF) token cookie and request parameters sent by the client exist and match. This is to prevent CSRF attacks.	true
<code>enable_gm_compiler</code>	boolean	If true, enable dynamic compilation of PGX Algorithm API (or Green-Marl code) during runtime.	true
<code>enable_shutdown_cleanup_hook</code>	boolean	If true, PGX will add a JVM shutdown hook that will automatically shutdown PGX at JVM shutdown. Notice: Having the shutdown hook deactivated and not explicitly shutting down PGX may result in pollution of your temp directory.	true
<code>enterprise_scheduler_config</code>	object	Configuration parameters for the enterprise scheduler.	null
<code>enterprise_scheduler_flags</code>	object	<i>[relevant for enterprise_scheduler]</i> Enterprise scheduler-specific settings.	null
<code>explicit_spin_locks</code>	boolean	true means spin explicitly in a loop until lock becomes available. false means using JDK locks which rely on the JVM to decide whether to context switch or spin. Setting this value to true usually results in better performance.	true
<code>file_locations</code>	array of objects	The file locations that can be used in the authorization-config.	[]
<code>graph_algorithm_language</code>	enum [GML, EGA, CY, GM, JAVA]	Front-end compiler to use.	gm

Table 11-1 (Cont.) Configuration Parameters for the Graph Server (PGX) Engine

Parameter	Type	Description	Default
graph_validation_level	enum[low, high]	Level of validation performed on newly loaded or created graphs.	low
ignore_incompatible_backend_operations	boolean	If true, only log when encountering incompatible operations and configuration values in RTS or FJ pool. If false, throw exceptions.	false
in_place_update_consistency_model	enum[ALL, ON_INCONSISTENCIES, CAN, CEL, _TA, SKS]	Consistency model used when in-place updates occur. Only relevant if in-place updates are enabled. Currently updates are only applied in place if the updates are not structural (Only modifies properties). Two models are currently implemented, one only delays new tasks when an update occurs, the other also delays running tasks.	all
init_pgql_on_startup	boolean	If true PGQL is directly initialized on start-up of PGX. Otherwise, it is initialized during the first use of PGQL.	true
interval_to_poll_max	integer	Exponential backoff upper bound (in ms) to which -once reached, the job status polling interval is fixed	100 0
java_home_dir	string	The path to Java's home directory. If set to <system-java-home-dir>, use the java.home system property.	null
large_array_threshold	integer	Threshold when the size of an array is too big to use a normal Java array. This depends on the used JVM. (Defaults to Integer.MAX_VALUE - 3)	214 748 364 4
max_active_sessions	integer	Maximum number of sessions allowed to be active at a time.	102 4
max_distinct_strings_per_pool	integer	<i>[only relevant if string_pooling_strategy is indexed]</i> Number of distinct strings per property after which to stop pooling. If the limit is reached, an exception is thrown.	655 36
max_http_client_request_size	long	Maximum size in bytes of any http request sent to the PGX server over the REST API. Setting it to -1 allows requests of any size.	104 857 60

**Table 11-1 (Cont.) Configuration Parameters for the Graph Server (PGX) Engine**


Parameter	Type	Description	Default
max_off_heap_size	integer	Maximum amount of off-heap memory (in megabytes) that PGX is allowed to allocate before an OutOfMemoryError will be thrown.	<available-physical-memory>
<div style="border: 1px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> <b>Note:</b></p> <p>This limit is not guaranteed to never be exceeded, because of rounding and synchronization trade-offs. It only serves as threshold when PGX starts to reject new memory allocation requests.</p> </div>			
max_queue_size_per_session	integer	The maximum number of pending tasks allowed to be in the queue, per session. If a session reaches the maximum, new incoming requests of that session get rejected. A negative value means infinity or unlimited..	-1
max_snapshot_count	integer	Number of snapshots that may be loaded in the engine at the same time. New snapshots can be created via auto or forced update. If the number of snapshots of a graph reaches this threshold, no more auto-updates will be performed, and a forced update will result in an exception until one or more snapshots are removed from memory. A value of zero indicates to support an unlimited amount of snapshots.	0
memory_allocator	enum[basicallocator, enterprise_allocator]	The memory allocator to use.	basic_allocator
memory_cleanup_interval	integer	Memory cleanup interval in seconds.	600

Table 11-1 (Cont.) Configuration Parameters for the Graph Server (PGX) Engine

Parameter	Type	Description	Default
<code>min_array_compaction_threshold</code>	number	Minimum value ( <i>only relevant for graphs optimized for updates</i> ) that can be used for the <code>array_compaction_threshold</code> value in graph configuration. If a graph configuration attempts to use a value lower than the one specified by <code>min_array_compaction_threshold</code> , it will use <code>min_array_compaction_threshold</code> instead.	0.2
<code>min_fetch_interval_sec</code>	integer	For delta-refresh ( <i>only relevant if the graph format supports delta updates</i> ), the lowest interval at which a graph source is queried for changes. You can tune this value to prevent PGX from hanging due to too frequent graph delta-refreshing.	2
<code>min_update_interval_sec</code>	integer	For auto-refresh, the lowest interval after which a new snapshot is created, either by reloading the entire graph or if the format supports delta-updates, out of the cached changes ( <i>only relevant if the format supports delta updates</i> ). You can tune this value to prevent PGX from hanging due to too frequent graph auto-refreshing.	2
<code>ms_bfs_frontier_type_strategy</code>	enum[auto_graph_show, int]	The type strategy to use for MS-BFS frontiers.	auto_graph
<code>num_spin_locks</code>	integer	Number of spin locks each generated app will create at instantiation. Trade-off: a small number implies less memory consumption; a large number implies faster execution (if algorithm uses spin locks).	1024
<code>parallelism</code>	integer	Number of worker threads to be used in thread pool. Note: If the caller thread is part of another thread-pool, this value is ignored and the parallelism of the parent pool is used.	<number of CPUs>
<code>pattern_matching_supernode_cache_threshold</code>	integer	Minimum number of a node's neighbor to be a supernode. This is for the pattern matching engine.	1000
<code>pgx_realm</code>	object	Configuration parameters for the realm.	null
<code>pooling_factor</code>	number	<i>[only relevant if string_pooling_strategy is on_heap]</i> This value prevents the string pool to grow as big as the property size, which could render the pooling ineffective.	0.25
<code>preload_graphs</code>	array of object	List of graph configs to be registered at start-up. Each item includes path to a graph config, the name of the graph and whether it should be published.	[]

Table 11-1 (Cont.) Configuration Parameters for the Graph Server (PGX) Engine

Parameter	Type	Description	Default
random_generator_strategy	enum	Method of generating random numbers in PGX.	nondeterministic
random_seed	long	<i>[relevant for deterministic random number generator only]</i> Seed for the deterministic random number generator used in pgx. The default is -24466691093057031.	-24466691093057031
release_memory_threshold	number	Threshold percentage (decimal fraction) of used memory after which the engine starts freeing unused graphs. Examples: A value of 0.0 means graphs get freed as soon as their reference count becomes zero. That is, all sessions which loaded that graph were destroyed/timed out. A value of 1.0 means graphs never get freed, and the engine will throw OutOfMemoryErrors as soon as a graph is needed which does not fit in memory anymore. A value of 0.7 means the engine keeps all graphs in memory as long as total memory consumption is below 70% of total available memory, even if there is currently no session using them. When consumption exceeds 70% and another graph needs to get loaded, unused graphs get freed until memory consumption is below 70% again.	0.85
revisit_threshold	integer	Maximum number of matched results from a node to be cached.	4096

**Table 11-1 (Cont.) Configuration Parameters for the Graph Server (PGX) Engine**

Parameter	Type	Description	Default
scheduler	enum[basic_scheduler, enterprise_scheduler, low_latency_scheduler]	The scheduler to use. <ul style="list-style-type: none"> <li>basic_scheduler: uses a scheduler with basic features</li> <li>enterprise_scheduler: uses a scheduler with advanced enterprise features for running multiple tasks concurrently and providing better performance</li> <li>low_latency_scheduler: uses a scheduler that privileges latency of tasks over throughput or fairness across multiple sessions. The low_latency_scheduler is only available in embedded mode.</li> </ul>	enterprise_scheduler
session_idle_timeout_secs	integer	Timeout of idling sessions in seconds. Zero (0) means infinity or no timeout	0
session_task_timeout_secs	integer	Timeout in seconds to interrupt long-running tasks submitted by sessions (algorithms, I/O tasks). Zero (0) means infinity or no timeout.	0
small_task_length	integer	Task length if the total amount of work is smaller than default task length (only relevant for task-stealing strategies).	128
strict_mode	boolean	If true, exceptions are thrown and logged with ERROR level whenever the engine encounters configuration problems, such as invalid keys, mismatches, and other potential errors. If false, the engine logs problems with ERROR/WARN level (depending on severity) and makes best guesses and uses sensible defaults instead of throwing exceptions.	true
string_pooling_strategy	enum[on_heap, non_heap]	The string pooling strategy to use.	on_heap



**Table 11-1 (Cont.) Configuration Parameters for the Graph Server (PGX) Engine**

Parameter	Type	Description	Default
task_length	integer	Default task length (only relevant for task-stealing strategies). Should be between 100 and 10000. Trade-off: a small number implies more fine-grained tasks are generated, higher stealing throughput; a large number implies less memory consumption and GC activity.	4096
tmp_dir	string	Temporary directory to store compilation artifacts and other temporary data. If set to <system-tmp-dir>, uses the standard tmp directory of the underlying system (/tmp on Linux).	null
udf_config_directory	string	Directory path containing UDF config files.	null
use_index_for_reachability_queries	enum[auto, off]	Create index for reachability queries.	auto
use_memory_mapper_for_reading_pgb	boolean	If true, use memory mapped files for reading graphs in PGB format if possible; if false, always use a stream-based implementation.	true
use_memory_mapper_for_storing_pgb	boolean	If true, use memory mapped files for storing graphs in PGB format if possible; if false, always use a stream-based implementation.	true

### Enterprise Scheduler Parameters

The following parameters are relevant only if the advanced scheduler is used. (They are ignored if the basic scheduler is used.)

- analysis\_task\_config  
Configuration for analysis tasks. Type: object. Default: prioritymediummax\_threads<no-of-CPU>weight<no-of-CPU>
- fast\_analysis\_task\_config  
Configuration for fast analysis tasks. Type: object. Default: priorityhighmax\_threads<no-of-CPU>weight1
- maxnum\_concurrent\_io\_tasks  
Maximum number of concurrent tasks. Type: integer. Default: 3
- num\_io\_threads\_per\_task  
Configuration for fast analysis tasks. Type: object. Default: <no-of-cpus>

### Basic Scheduler Parameters

The following parameters are relevant only if the basic scheduler is used. (They are ignored if the advanced scheduler is used.)

- num\_workers\_analysis

Number of worker threads to use for analysis tasks. Type: integer. Default: <no-of-CPUs>

- `num_workers_fast_track_analysis`

Number of worker threads to use for fast-track analysis tasks. Type: integer. Default: 1

- `num_workers_io`

Number of worker threads to use for I/O tasks (load/refresh/write from/to disk). This value will not affect file-based loaders, because they are always single-threaded. Database loaders will open a new connection for each I/O worker. Default: <no-of-CPUs>

### Example 11-1 Minimal In-Memory Graph Server (PGX) Configuration

The following example causes the in-memory graph server (PGX) to initialize its analysis thread pool with 32 workers. (Default values are used for all other parameters.)

```
{
 "enterprise_scheduler_config": {
 "analysis_task_config": {
 "max_threads": 32
 }
 }
}
```

### Example 11-2 Two Pre-loaded Graphs

This example sets more fields and specifies two fixed graphs for loading into memory during the graph server (PGX) startup.

```
{
 "enterprise_scheduler_config": {
 "analysis_task_config": {
 "max_threads": 32
 },
 "fast_analysis_task_config": {
 "max_threads": 32
 }
 },
 "memory_cleanup_interval": 600,
 "max_active_sessions": 1,
 "release_memory_threshold": 0.2,
 "preload_graphs": [
 {
 "path": "graph-configs/my-graph.bin.json",
 "name": "my-graph"
 },
 {
 "path": "graph-configs/my-other-graph.adj.json",
 "name": "my-other-graph",
 "publish": false
 }
],
 "authorization": [{
 "pgx_role": "GRAPH_DEVELOPER",
 "pgx_permissions": [{
 "preloaded_graph": "my-graph",
 "grant": "read"
 }
]
}]
}
```

```

 },
 {
 "preloaded_graph": "my-other-graph",
 "grant": "read"
 }
],

]
}

```

- [Configuration of the Graph Server \(PGX\) Run-Time Parameters](#)
- [Specifying the Configuration File to the In-Memory Graph Server \(PGX\)](#)
- [Memory Consumption by the Graph Server \(PGX\)](#)  
The in-memory graph server (PGX) loads the graph into main memory in order to carry out analysis on the graph and its properties.

## 11.1.1 Configuration of the Graph Server (PGX) Run-Time Parameters

You can configure the following graph server (PGX) run-time fields.

**Table 11-2 Graph Server (PGX) Run-Time Parameters**

Parameter	Type	Description	Default
bfs_iterate_que_task_size	integer	Task size for BFS iterate QUE phase.	128
bfs_threshold_parent_read_based	number	Threshold of BFS traversal level items above which to switch to parent-read-based visiting strategy.	0.05
bfs_threshold_read_based	integer	Threshold of BFS traversal level items above which to switch to read-based visiting strategy.	1024
bfs_threshold_single_threaded	integer	Number until which BFS traversal level items vertices are visited single-threaded.	128
character_set	string	Standard charset to use throughout PGX, UTF-8 will be used as default. Note: Some formats may not be compatible.	utf-8
cni_diff_factor_default	integer	Default diff factor value used in the common neighbor iterator implementations.	8
cni_small_default	integer	Default value used in the common neighbor iterator implementations, to indicate below which threshold a subarray is considered small.	128
cni_stop_recursion_default	integer	Default value used in the common neighbor iterator implementations, to indicate the minimum size where the binary search approach is applied.	96
dfs_threshold_large	integer	Value that determines at which number of visited vertices, the DFS implementation will switch to data-structures that are more optimized for larger numbers of vertices.	4096
enterprise_scheduler_flags	object	<i>[relevant for enterprise_scheduler]</i> Enterprise scheduler specific settings.	null

**Table 11-2 (Cont.) Graph Server (PGX) Run-Time Parameters**


Parameter	Type	Description	Default
<code>explicit_spin_locks</code>	boolean	<code>true</code> means spin explicitly in a loop until lock becomes available. <code>false</code> means using JDK locks which rely on the JVM to decide whether to context switch or spin. Our experiments showed that setting this value to <code>true</code> results in better performance.	<code>true</code>
<code>graph_validation_level</code>	enum[low, high]	Level of validation performed on newly loaded or created graphs.	low
<code>max_distinct_strings_per_pool</code>	integer	<i>[only relevant if string pooling strategy is indexed]</i> Amount of distinct strings per property after which to stop pooling. If the limit is reached an exception is thrown.	65536
<code>max_off_heap_size</code>	integer	Maximum amount of off-heap memory PGX is allowed to allocate in megabytes, before an <code>OutOfMemoryError</code> will be thrown.	<available-physical-memory>
<div style="border: 1px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> <b>Note:</b></p> <p>This limit is not guaranteed to never be exceeded because of rounding and synchronization trade-offs. It only serves as threshold when PGX starts to reject new memory allocation requests.</p> </div>			
<code>memory_allocator</code>	enum[basic_allocator, enterprise_allocator]	Denotes which memory allocator to use.	basic_allocator
<code>ms_bfs_frontier_type_strategy</code>	enum[auto_grow, short, int]	The type strategy to use for MS-BFS frontiers.	auto_grow
<code>num_spin_locks</code>	integer	Number of spin locks each generated app will create at instantiation. Trade-off: small number implies less memory consumption. Big number implies faster execution (if algorithm uses spin locks).	1024
<code>pattern_matching_supernode_cache_threshold</code>	integer	Minimum number of a node's neighbor to be a supernode. This is for pattern matching engine.	1000

Table 11-2 (Cont.) Graph Server (PGX) Run-Time Parameters

Parameter	Type	Description	Default
pooling_factor	number	<i>[only relevant if string_pooling_strategy is on_heap]</i> This value prevents the string pool to grow as big as the property size which could render the pooling ineffective.	0.25
random_generator_strategy	enum[non_deterministic, deterministic]	Method of generating random numbers in PGX.	non_deterministic
random_seed	long	<i>[relevant for deterministic random number generator only]</i> Seed for the deterministic random number generator used in PGX. The default is -24466691093057031.	-244666 9109305 7031
revisit_threshold	integer	Maximum number of matched results from a node to be cached.	4096
scheduler	enum[basic_scheduler, enterprise_scheduler, low_latency_scheduler]	Denotes which scheduler to use. <ul style="list-style-type: none"> <li>basic_scheduler: use scheduler with basic features.</li> <li>enterprise_scheduler: use scheduler with advanced, enterprise features for running multiple tasks concurrently and increased performance.</li> <li>low_latency_scheduler: use scheduler that privileges latency of tasks over throughput or fairness across multiple sessions. The low_latency_scheduler is only available in embedded mode</li> </ul>	enterprise_scheduler
small_task_length	integer	Task length, if total amount of work is small than default task length (only relevant for task-stealing strategies).	128
string_pooling_strategy	enum[in_dexed, on_heap, none]	Denotes which string pooling strategy to use.	on_heap
task_length	integer	Default task length (only relevant for task-stealing strategies). F/J pool documentation says this value should be between 100 and 10000. Trade-off: small number implies more fine-grained tasks are generated, higher stealing throughput. High number implies less memory consumption and GC activity.	4096
use_index_for_reachability_queries	enum[auto, off]	Create index for reachability queries.	auto
use_memory_mapper_for_reading_pgb	boolean	If true, use memory mapped files for reading graphs in PGB format if possible; false always use stream based implementation.	true

**Table 11-2 (Cont.) Graph Server (PGX) Run-Time Parameters**

Parameter	Type	Description	Default
<code>use_memory_mapper_for_storing_pgb</code>	boolean	If <code>true</code> , use memory mapped files for storing in PGB format if possible; if <code>false</code> always use a stream based implementation.	<code>true</code>

## 11.1.2 Specifying the Configuration File to the In-Memory Graph Server (PGX)

The in-memory graph server configuration file is parsed by the in-memory graph server at startup-time whenever `ServerInstance#startEngine` (or any of its variants) is called. You can write the path to your configuration file to the in-memory graph server or specify it programmatically. This topic identifies several ways to specify the file

### Programmatically

All configuration fields exist as Java enums. Example:

```
Map<PgxCfg.Field, Object> pgxCfg = new HashMap<>();
pgxCfg.put(PgxCfg.Field.MEMORY_CLEANUP_INTERVAL, 600);

ServerInstance instance = ...
instance.startEngine(pgxCfg);
```

All parameters not explicitly set will get default values.

### Explicitly Using a File

Instead of a map, you can write the path to an in-memory graph server configuration JSON file. Example:

```
instance.startEngine("path/to/pgx.conf"); // file on local file system
instance.startEngine("classpath:/path/to/pgx.conf"); // file on current classpath
```

For all other protocols, you can write directly in the input stream to a JSON file. Example:

```
InputStream is = ...
instance.startEngine(is);
```

### Implicitly Using a File

If `startEngine()` is called without an argument, the in-memory graph server (PGX) looks for a configuration file at the following places, stopping when it finds the file:

- File path found in the Java system property `pgx_conf`. Example: `java -Dpgx_conf=conf/my.pgx.config.json ...`
- A file named `pgx.conf` in the root directory of the current classpath
- A file named `pgx.conf` in the root directory relative to the current `System.getProperty("user.dir")` directory

Note: Providing a configuration is optional. A default value for each field will be used if the field cannot be found in the given configuration file, or if no configuration file is provided.

### Using the Shell in Embedded Mode

To change how the shell configures the embedded (local) in-memory graph server (PGX) instance, edit `$PGX_HOME/conf/pgx.conf`. Changes will be reflected the next time you invoke `$PGX_HOME/bin/pgx`.

You can also change the location of the configuration file as in the following example:

```
./bin/opg --pgx_conf path/to/my/other/pgx.conf
```

### Setting System Properties

Any parameter can be set using Java system properties by writing `-Dpgx.<FIELD>=<VALUE>` arguments to the JVM that the in-memory graph server (PGX) is running on. Note that setting system properties will overwrite any other configuration. The following example sets the maximum off-heap size to 256 GB, regardless of what any other configuration says:

```
java -Dpgx.max_off_heap_size=256000 ...
```

### Setting Environment Variables

Any parameter can also be set using environment variables by adding 'PGX\_' to the environment variable for the JVM in which the in-memory graph server (PGX) is executed. Note that setting environment variables will overwrite any other configuration; but if a system property and an environment variable are set for the same parameter, the system property value is used. The following example sets the maximum off-heap size to 256 GB using an environment variable:

```
PGX_MAX_OFF_HEAP_SIZE=256000 java ...
```

## 11.1.3 Memory Consumption by the Graph Server (PGX)

The in-memory graph server (PGX) loads the graph into main memory in order to carry out analysis on the graph and its properties.

The memory consumed by the graph server for a graph is split between the memory to store the topology of the graph (the information to indicate what are the vertices and edges in the graph without their attached properties), and the memory for the properties attached to the vertices and edges. Internally, the graph server (PGX) stores the graph topology in compressed sparse row (CSR) format, a data structure which has minimal memory footprint while providing very fast read access.

- [Memory Management](#)

### 11.1.3.1 Memory Management

The in-memory graph server (PGX) requires both on-heap and off-heap memory to store graph data.

The allocation of memory for the graph data is as shown:

- Graph indexes and graph topology are stored off-heap.
- All primitive properties (integer, long, double, float, boolean, date, local\_date, timestamp, time, point2d) are stored off-heap.

- String properties are stored on-heap.

### Default Configuration of Memory Limits

You can configure both on-heap and off-heap memory limits. In case of the on-heap, if you don't explicitly set a maximum then it will default to the maximum on-heap size determined by Java Hotspot, which is based on various factors, including the total amount of physical memory available. In case of the off-heap, if you don't explicitly set a maximum then it will default to the total physical available memory on the machine.

- [Configuring On-Heap Limits](#)
- [Configuring Off-Heap Limits](#)

#### 11.1.3.1.1 Configuring On-Heap Limits

You can configure on-heap limits using Java command-line options.

The available options are:

- `-Xmx`: to set the maximum on-heap size of the JVM.
- `-Xms`: to set the initial on-heap size of the JVM.
- `-XX:NewSize`: to set the initial size of the young generation
- `-XX:MaxNewSize`: to set the maximum size of the young generation

The following shows an example to configure the on-heap limits using `-XX:MaxNewSize` option in a Java application:

```
java -Xmx<size_mb>m -Xms<size_mb>m -XX:MaxNewSize=<size_mb>m -
XX:NewSize=<size_mb>m
```

If you are using a JShell client, then you can set the `JAVA_OPTS` environment variable before starting the shell. For example:

```
export JAVA_OPTS="-Xmx<size_gb>g -Xms<size_gb>g -
XX:MaxNewSize=<size_gb>g -XX:NewSize=<size_gb>g "
cd /opt/oracle/graph/
./bin/opg-jshell
```

#### 11.1.3.1.2 Configuring Off-Heap Limits

You can specify the off-heap limit by setting the `max_off_heap_size` field in the graph server (PGX) configuration. See [Configuration Parameters for the Graph Server \(PGX\) Engine](#) for more information on the `max_off_heap_size` parameter.

#### WARNING:

The off-heap limit is not guaranteed to never be exceeded because of rounding and synchronization trade-offs.

The off-heap limit can be set using Java system properties using `-Dpgx.max_off_heap_size=<size_in_mb>` in the JVM argument.



You can also set the off-heap limit using an environment variable. The following example sets the maximum off-heap size to 256 GB using an environment variable:

```
PGX_MAX_OFF_HEAP_SIZE=256000 java ...
```



#### Note:

If both system property and environment variable are set for off-heap limit, then the system property value is used.

## 11.2 Configuration Parameters for Connecting to the Graph Server (PGX)

You can configure the graph server (PGX) to use the required options at startup.

See [Configuring the In-Memory Graph Server \(PGX\)](#)

## 11.3 Configuration Parameters for the Graph Client

You can configure the PGX graph client. All the parameters are available as command-line options also.

**Table 11-3 Configuration Parameters for the Graph Client**

Parameter	Type	Description	Default
<code>access_token</code>	string	The authentication token.	null
<code>base_url</code>	string	The base url in the format <code>host [ : port] [ /path]</code> of the PGX server REST endpoint. If the <code>base_url</code> is null, the default will be used which points to embedded PGX instance.	null
<code>cctrace_out</code>	string	<i>[relevant for <code>enable_cctrace</code>]</i> When <code>cctrace</code> is enabled, this option specifies a path to a file where <code>cctrace</code> should log to. If null it will use the default PGX logger on level TRACE. If it is the special value <code>:stderr</code> : it will log to <code>stderr</code> .	null
<code>cctrace_print_stacktraces</code>	boolean	<i>[relevant for <code>enable_cctrace</code>]</i> When <code>cctrace</code> is enabled, this flag prints the stacktrace for each request and result.	false
<code>client_server_interaction_mode</code>	enum[ <code>async_polling</code> , <code>blocking</code> ]	If <code>async_polling</code> the PGX client would poll the status of the future until it is completed. If <code>blocking</code> , the PGX client would send a request to directly get the value of the future and the server would block until the future result is ready.	<code>async_polling</code>
<code>enable_cctrace</code>	boolean	If true log every call to a Control or Core interface.	false

**Table 11-3 (Cont.) Configuration Parameters for the Graph Client**

Parameter	Type	Description	Default
keystore	string	The path to the keystore to use for client connections. The keystore is used to authenticate this client at the PGX server if two-way SSL/TLS is enabled.	null
max_client_http_connections	integer	Maximum number of connections to open to the PGX server.	2
password	string	Keystore password only.	null
prefetch_size	integer	Number of items to be prefetched in remote iterators.	2048
realm_client_config	object	Implementation dependent configuration options for the realm client.	null
remote_future_pending_retry_interval	integer	Number of milliseconds to wait before sending another request in case a GET request for a <code>PgxRemoteFuture</code> receives a 202 - Accepted response.	500
remote_future_timeout	integer	Time that a GET request for a <code>PgxRemoteFuture</code> will be alive, until it times out and tries again. Time in milliseconds, set it to zero for an infinite timeout. See HTTP Client <code>SO_TIMEOUT</code> for more details.	300000
tls_version	string	TLS version to be used by the client. For example, <code>TLSv1.2</code> .	<code>tlsv1.2</code>
truststore	string	Path to the truststore to use for client connections. The truststore is used to validate the server certificate if communicating over SSL/TLS.	null
upload_batch_size	integer	Number of items to be uploaded in a batch. This is used in <code>Core#addAllToCollection()</code> and <code>Core#setProperty()</code> .	65536
username	string	Name of the user.	null

**Example 11-3 Configure the Graph Client Using the Graph PGX Shell**

This following is an example to configure the graph client:

```
cd /opt/oracle/graph
./bin/opg-jshell --base_url https://myhost:8080/pgx --username scott --
prefetch_size 1024 --upload_batch_size 5000 --remote_future_timeout 20000 --
pending_retry_interval 800
```

**Example 11-4 Configure the Graph Client Using the Java API**

The following is an example to configure the graph client programmatically using the `Pgx.getInstance` methods:

```
public static ServerInstance getInstance(String baseUrl, String
username, String password, Integer prefetchSize,
```

```
Integer uploadBatchSize, Integer remoteFutureTimeout, Integer
remoteFuturePendingRetryInterval)
```

To specify key store and trust store for SSL connections use the standard JDK system properties:

```
System.setProperty("javax.net.ssl.trustStore", "<truststore>");
System.setProperty("javax.net.ssl.keyStore", "<keystore>");
System.setProperty("javax.net.ssl.keyStorePassword", "<password>");
```

# 12

## Graphs Management

You can load, publish, store and delete graphs

- [Loading a Graph Into the Graph Server \(PGX\)](#)
- [Publishing a Graph](#)
- [Publishing a Preloaded Graph](#)
- [Deleting a Graph](#)

### 12.1 Loading a Graph Into the Graph Server (PGX)

Data from relational tables can be modeled as a property graph and loaded into the graph server.

The graph server (PGX) supports various data sources and data formats for loading graph data, including file system and database formats. See [Data Format Support Matrix](#) to get more information on the supported data formats.

In order to perform graph analysis with the graph server (PGX), you must first read a graph into PGX. See [Reading Graphs from Oracle Database into the Graph Server \(PGX\)](#) for details.

- [API for Loading Graphs into Memory](#)
- [Graph Configuration Options](#)
- [Preloading a Graph](#)
- [Data Loading Security Best Practices](#)
- [Data Format Support Matrix](#)
- [Immutability of Loaded Graphs](#)

#### 12.1.1 API for Loading Graphs into Memory

The following methods in `PgxSession` can be used to load graphs into the graph server (PGX) memory:

- 
- [Java](#)
  - [Python](#)

#### Java

```
PgxGraph readGraphWithProperties(String path)
PgxGraph readGraphWithProperties(String path, String newGraphName)
PgxGraph readGraphWithProperties(GraphConfig config)
```

```
PgxGraph readGraphWithProperties(GraphConfig config, String
newGraphName)
PgxGraph readGraphWithProperties(GraphConfig config, boolean
forceUpdateIfNotFresh)
PgxGraph readGraphWithProperties(GraphConfig config, boolean
forceUpdateIfNotFresh, String newGraphName)
PgxGraph readGraphWithProperties(GraphConfig config, long maxAge,
TimeUnit maxAgeTimeUnit)
PgxGraph readGraphWithProperties(GraphConfig config, long maxAge,
TimeUnit maxAgeTimeUnit, boolean blockIfFull, String newGraphName)
```

## Python

```
read_graph_with_properties(self, config, max_age=9223372036854775807,
max_age_time_unit='days',
 block_if_full=False,
update_if_not_fresh=True, graph_name=None)
```

The first argument (path to a graph config file or a parsed `config` object) is the meta-data of the graph to be read. The meta-data includes the following information:

- Location of the graph data: file location and name, DB location and connection information and so on
- Format of the graph data: plain text formats, XML-based formats, Binary formats and so on
- Types and Names of the properties to be loaded

The `forceUpdateIfNotFresh` and `maxAge` arguments can be used to fine-control the age of the snapshot to be read. The graph server (PGX) will return an existing graph snapshot if the given graph specification was already loaded into memory by a different session. So, the `maxAge` argument becomes important if reading from a database in which the data might change frequently. If no `forceUpdateIfNotFresh` or `maxAge` is specified, PGX will favor cached data over reading new snapshots into memory.

### 12.1.2 Graph Configuration Options

The following table lists the JSON fields that are common to all graph configurations:

**Table 12-1 Graph Config JSON Fields**

Field	Type	Description	Default
name	string	Name of the graph.	Required

Table 12-1 (Cont.) Graph Config JSON Fields

Field	Type	Description	Default
array_compaction_threshold	number	<i>[only relevant if the graph is optimized for updates]</i> Threshold used to determine when to compact the delta-logs into a new array. If lower than the engine <code>min_array_compaction_threshold</code> value, <code>min_array_compaction_threshold</code> will be used instead	0.2
attributes	object	Additional attributes needed to read and write the graph data.	null
edge_id_strategy	enum[no_ids, keys_as_ids, unstable_generated_ids]	Indicates what ID strategy should be used for the edges of this graph. If not specified (or set to null), the strategy will be determined during loading or using a default value.	null
edge_id_type	enum[long]	Type of the edge ID. Setting it to <code>long</code> requires the IDs in the edge providers to be unique across the graphs; those IDs will be used as global IDs. Setting it to <code>null</code> (or omitting it) will allow repeated IDs across different edge providers and PGX will automatically generate globally-unique IDs for the edges.	null
edge_providers	array of object	List of edge providers in this graph.	[]
error_handling	object	Error handling configuration.	null
external_stores	array of object	Specification of the external stores where external string properties reside.	[]
jdbc_url	string	JDBC URL pointing to an RDBMS instance	null
keystore_alias	string	Alias to the keystore to use when connecting to database.	null
loading	object	Loading-specific configuration to use.	null
local_date_formats	array of string	array of <code>local_date</code> formats to use when loading and storing <code>local_date</code> properties. See <a href="#">DateTimeFormatter</a> for more details of the format string	[]
max_prefetched_rows	integer	Maximum number of rows prefetched during each round trip resultset-database.	10000
num_connections	integer	Number of connections to read and write data from or to the RDBMS table.	<no-of-cpus>
optimized_for	enum[read, updates]	Indicates if the graph should use data-structures optimized for read-intensive scenarios or for fast updates.	read
password	string	Password to use when connecting to database.	null

**Table 12-1 (Cont.) Graph Config JSON Fields**

Field	Type	Description	Default
point2d	string	Longitude and latitude as floating point values separated by a space.	0.0 0.0
redaction_rules	array of object	Array of redaction rules.	[]
rules_mapping	array of object	Mapping for redaction rules to users and roles.	[]
schema	string	Schema to use when reading or writing RDBMS objects	null
time_format	array of string	The time format to use when loading and storing time properties. See <a href="#">DateTimeFormatter</a> for a documentation of the format string.	[]
time_with_timezone_one_format	array of string	The time with timezone format to use when loading and storing time with timezone properties. Please see <a href="#">DateTimeFormatter</a> for more information of the format string.	[]
timestamp_format	array of string	The timestamp format to use when loading and storing timestamp properties. See <a href="#">DateTimeFormatter</a> for more information of the format string.	[]
timestamp_with_timezone_format	array of string	The timestamp with timezone format to use when loading and storing timestamp with timezone properties. See <a href="#">DateTimeFormatter</a> for more information of the format string.	[]
username	string	Username to use when connecting to an RDBMS instance.	null
vector_component_delimiter	character	Delimiter for the different components of vector properties.	;
vertex_id_strategy	enum[no_ids, keys_as_ids, unstable_generated_ids]	Indicates what ID strategy should be used for the vertices of this graph. If not specified (or set to null), the strategy will be automatically detected.	null
vertex_id_type	enum[int, integer, long, string]	Type of the vertex ID. For homogeneous graphs, if not specified (or set to null), it will default to a specific value (depending on the origin of the data).	null
vertex_providers	array of object	List of vertex providers in this graph.	[]

**Note:**

Database connection fields specified in the graph configuration will be used as default in case underlying data provider configuration does not specify them.

**Provider Configuration JSON file Options**

You can specify the meta-information about each provider's data using provider configurations. Provider configurations include the following information about the provider data:

- Location of the data: a file, multiple files or database providers
- Information about the properties: name and type of the property

**Table 12-2 Provider Configuration JSON file Options**

Field	Type	Description	Default
format	enum[pgb, csv, rdbms]	Provider format.	Required
name	string	Entity provider name.	Required
attributes	object	Additional attributes needed to read and write the graph data.	null
destination_vertex_provider	string	Name of the destination vertex provider to be used for this edge provider.	null
error_handling	object	Error handling configuration.	null
has_keys	boolean	Indicates if the provided entities data have keys.	true
key_type	enum[int, integer, long, string]	Type of the keys.	long
keystore_alias	string	Alias to the keystore to use when connecting to database.	null
label	string	label for the entities loaded from this provider.	null
loading	object	Loading-specific configuration.	null
local_date_formats	array of string	Array of local_date formats to use when loading and storing local_date properties. See <a href="#">DateTimeFormatter</a> for a documentation of the format string.	[]
password	string	Password to use when connecting to database.	null
point2d	string	Longitude and latitude as floating point values separated by a space.	0.0 0.0
props	array of object	Specification of the properties associated with this entity provider.	[]
source_vertex_provider	string	Name of the source vertex provider to be used for this edge provider.	null



**Table 12-2 (Cont.) Provider Configuration JSON file Options**

Field	Type	Description	Default
<code>time_format</code>	array of string	The time format to use when loading and storing time properties. See <a href="#">DateTimeFormatter</a> for a documentation of the format string.	[]
<code>time_with_timezone_format</code>	array of string	The time with timezone format to use when loading and storing time with timezone properties. See <a href="#">DateTimeFormatter</a> for a documentation of the format string.	[]
<code>timestamp_format</code>	array of string	The timestamp format to use when loading and storing timestamp properties. See <a href="#">DateTimeFormatter</a> for a documentation of the format string.	[]
<code>timestamp_with_timezone_format</code>	array of string	The timestamp with timezone format to use when loading and storing timestamp with timezone properties. See <a href="#">DateTimeFormatter</a> for a documentation of the format string.	[]
<code>vector_component_delimiter</code>	character	Delimiter for the different components of vector properties.	;

### Provider Labels

The `label` field in the provider configuration can be used to set a label for the entities loaded from the provider. If no `label` is specified, all entities from the provider are labeled with the name of the provider. It is only possible to set the same label for two different providers if they have exactly the same properties (same names and same types).

### Property Configuration

The `props` entry in the `Provider` configuration is an object with the following JSON fields:

**Table 12-3 Property Configuration**

Field	Type	Description	Default
<code>name</code>	string	Name of the property.	Required

**Table 12-3 (Cont.) Property Configuration**

Field	Type	Description	Default
type	enum[boolean, integer, vertex, edge, float, long, double, string, date, local_date, time, timestamp, time_with_timezone, timestamp_with_timezone, point2d]	Type of the property .	Required
		vertex/edge are place-holders for the type specified in vertex_id_type/edge_id_type fields.	
aggregate	enum[identity, group_key, min, max, avg, sum, concat, count]	[currently unsupported] which aggregation function to use, aggregation always happens by vertex key.	null
column	value	Name or index (starting from 0) of the column holding the property data. If it is not specified, the loader will try to use the property name as column name (for CSV format only).	null
default	value	Default value to be assigned to this property if datasource does not provide it. In case of date type: string is expected to be formatted with yyyy-MM-dd HH:mm:ss. If no default is present (null), non-existent properties will contain default Java types (primitives) or empty string (string) or 01.01.1970 00:00 (date).	null
dimension	integer	Dimension of property.	0
drop_after_loading	boolean	[currently unsupported] indicating helper properties only used for aggregation, which are dropped after loading	false

 **Note:**

date is deprecated, use one of local\_date / time / timestamp / time\_with\_timezone / timestamp\_with\_timezone instead).

**Table 12-3 (Cont.) Property Configuration**

Field	Type	Description	Default
field	value	Name of the JSON field holding the property data. Nesting is denoted by dot - separation. Field names containing dots are possible, in this case the dots need to be escaped using backslashes to resolve ambiguities. Only the exactly specified object are loaded, if they are non existent, the default value is used.	null
format	array of string	Array of formats of property.	[]
group_key	string	[currently unsupported] can only be used if the property / key is part of the grouping expression.	null
max_distinct_strings_per_pool	integer	<i>[only relevant if string_pooling_strategy is indexed]</i> Amount of distinct strings per property after which to stop pooling. If the limit is reached an exception is thrown. If set to null, the default value from the global PGX configuration will be used.	null
stores	array of object	A list of storage identifiers that indicate where this property resides.	[]
string_pooling_strategy	enum[indexed, on_heap, none]	Indicates which string pooling strategy to use. If set to null, the default value from the global PGX configuration will be used.	null

### Loading Configuration

The loading entry is a JSON object with the following fields:

**Table 12-4 Loading Configuration**

Field	Type	Description	Default
create_key_mapping	boolean	If true, a mapping between entity keys and internal IDs is prepared during loading.	true
filter	string	[currently unsupported] the filter expression	null
grouping_by	array of string	[currently unsupported] array of edge properties used for aggregator. For Vertices, only the ID can be used (default)	[]
load_labels	boolean	Whether or not to load the entity label if it is available.	false

**Table 12-4 (Cont.) Loading Configuration**

Field	Type	Description	Default
<code>strict_mode</code>	boolean	If <code>true</code> , exceptions are thrown and logged with <code>ERROR</code> level whenever loader encounters problems with input file, such as invalid format, repeated keys, missing fields, mismatches and other potential errors. If <code>false</code> , loader may use less memory during loading phase, but behave unexpectedly with erratic input files.	<code>true</code>

**Error Handling Configuration**

The `error_handling` entry is a JSON object with the following fields:

**Table 12-5 Error Handling Configuration**

Field	Type	Description	Default
<code>on_missed_prop_key</code>	enum[ <code>silent</code> , <code>log_warn</code> , <code>log_warn_once</code> , <code>error</code> ]	Error handling for a missing property key.	<code>log_warn_once</code>
<code>on_missing_vertex</code>	enum[ <code>ignore_edge</code> , <code>ignore_edge_log</code> , <code>ignore_edge_log_once</code> , <code>create_vertex</code> , <code>create_vertex_log</code> , <code>create_vertex_log_once</code> , <code>error</code> ]	Error handling for a missing source or destination vertex of an edge in a vertex data source.	<code>error</code>
<code>on_parsing_issue</code>	enum[ <code>silent</code> , <code>log_warn</code> , <code>log_warn_once</code> , <code>error</code> ]	Error handling for incorrect data parsing. If set to <code>silent</code> , <code>log_warn</code> or <code>log_warn_once</code> , will attempt to continue loading. Some parsing issues may not be recoverable and provoke the end of loading.	<code>error</code>
<code>on_prop_conversion</code>	enum[ <code>silent</code> , <code>log_warn</code> , <code>log_warn_once</code> , <code>error</code> ]	Error handling when encountering a different property type other than the one specified, but coercion is possible.	<code>log_warn_once</code>
<code>on_type_mismatch</code>	enum[ <code>silent</code> , <code>log_warn</code> , <code>log_warn_once</code> , <code>error</code> ]	Error handling when encountering a different property type other than the one specified, but coercion is <i>not</i> possible.	<code>error</code>
<code>on_vector_length_mismatch</code>	enum[ <code>silent</code> , <code>log_warn</code> , <code>log_warn_once</code> , <code>error</code> ]	Error handling for a vector property that does not have the correct dimension.	<code>error</code>

**Note:**

The only supported setting for the `on_missing_vertex` error handling configuration is `ignore_edge`.

### 12.1.3 Preloading a Graph

You can configure the graph server (PGX) to preload graphs in memory at startup-time. This can be useful when you want the graph server (PGX) to startup automatically and have a graph (or multiple graphs) ready for its users. For example, deploying the graph server (PGX) on Kubernetes can be one such scenario.

The configuration for this is done through the `preload_graphs` configuration field in the graph server (PGX) configuration file.

The following is an example for preloading a graph using the graph configuration file:

```
{
 "preload_graphs": [
 {
 "path": "<path-to-graph-config>",
 "name": "my-graph"
 }
],
 "authorization": [{
 "pgx_role": "GRAPH_DEVELOPER",
 "pgx_permissions": [{
 "preloaded_graph": "my-graph",
 "grant": "read"
 }]
 }],

}
```

You can access a preloaded graph by its name using the `getGraph()` method of the session object.

```
PgxGraph g = session.getGraph("my-graph");
```

### 12.1.4 Data Loading Security Best Practices

Loading graph from the database requires authentication and it is therefore recommended to adhere to the following guidelines when configuring access to this kind of data source:

- The user or role used to access the data should be a read-only account that only has access to the required graph data.
- The graph data should be marked as read-only, for example, with non-updateable views in the case of the database.

## 12.1.5 Data Format Support Matrix

The following table illustrates how the different data formats differ in the way IDs, labels and vector properties are handled.



### Note:

The table refers to limitations of the PGX implementation of the format and not necessarily to limitations of the format itself.

**Table 12-6 Data Format Support Matrix**

Format	Vertex IDs	Edge IDs	Vertex Labels	Edge Labels	Vector properties
PGB	int, long, string	long	multiple	single	supported (vectors can be of type integer, long, float or double)
CSV	int, long, string	long	multiple	single	supported (vectors can be of type integer, long, float or double)
ADJ_LIST	int, long, string	not supported	not supported	not supported	supported (vectors can be of type integer, long, float or double)
EDGE_LIST	int, long, string	not supported	multiple	single	supported (vectors can be of type integer, long, float or double)
GRAPHML	int, long, string	not supported	not supported	not supported	not supported
TWO_TABLES	int, long, string	long	multiple	single	only in text datastore (vectors can be of type integer, long, float or double)
PG (FLAT_FILE)	int, long	long		single	not supported

## 12.1.6 Immutability of Loaded Graphs

The graph, once loaded into the graph server (PGX), the graph and its properties are automatically marked as immutable.

The immutability of loaded graphs is due to the following design choices:

- Typical graph analyses happen on a snapshot of a graph instance, and therefore they do not require mutations of the graph instance.
- Immutability allows PGX to use an internal graph representation optimized for fast analysis.
- In remote mode, the graph instance might be shared among multiple clients.

However, the graph server (PGX) also provides methods to privatize and mutate graph instances for the purpose of analysis. See [Graph Mutation and Subgraphs](#) for more information.

## 12.2 Publishing a Graph

### Publishing a Single Graph Snapshot

The `publish()` methods in `PgxGraph` can be used to publish the current selected snapshot of the graph.



#### Note:

Calling `publish()` without arguments publishes the snapshot with its persistent properties but does not publish transient properties.

This operation will move the graph name from the session-private namespace to the public namespace (see [Namespaces and Sharing](#) for more information about namespaces). If a graph with the same name has been already published, the `publish()` method will fail with an exception.



#### Note:

Graphs published with snapshots and single published snapshots share the same namespace.

For example, see [Example 4-1](#) to publish a graph using `publish()` method.

If you want to publish specific transient properties, you must list them within the `publish()` call.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
opg4j> var prop1 = graph.createVertexProperty(PropertyType.INTEGER,
"prop1")
opg4j> prop.fill(0)
opg4j> var cost = graph.createEdgeProperty(PropertyType.DOUBLE, "cost")
opg4j> cost.fill(0d)
opg4j> graph.publish(List.of(prop1), List.of(cost))
```

## Java

```
VertexProperty<Integer, Integer> prop1 =
graph.createVertexProperty(PropertyType.INTEGER, "prop1");
prop.fill(0);
EdgeProperty<Double> cost = graph.createEdgeProperty(PropertyType.DOUBLE,
"cost");
cost.fill(0d);
List<VertexProperty<Integer, Integer> vertexProps = Arrays.asList(prop);
List<EdgeProperty<Double>> edgeProps = Arrays.asList(cost);
graph.publish(vertexProps, edgeProps);
```

## Python

```
prop = graph.create_vertex_property("integer", "prop1")
prop.fill(0)
cost = graph.create_edge_property("double", "cost")
cost.fill(0d)
vertex_props = [prop]
edge_props = [cost]
graph.publish(vertex_props, edge_props)
```

---

### Publishing a Graph with Snapshots

If you want to make all snapshots of the graph visible to other sessions, use the `publishWithSnapshots()` methods instead. When a graph is published with snapshots, the `GraphMetaData` information of each snapshot is also made available to the other sessions, with the exception of the graph configuration, which is `null`.

With publishing, all persistent properties of all snapshots are also published and made visible to the other sessions, while transient properties are session-private and thus should be published explicitly. Once published, all properties become read-only. Hence, transient properties are not published when calling `publishWithSnapshots()` without arguments.

Similar to publishing a single graph snapshot, `publishWithSnapshots()` method will move the graph name from the session-private namespace to the public namespace (see [Namespaces and Sharing](#) for more information about namespaces). If a graph with the same name has been already published, the `publishWithSnapshots()` method will fail with an exception.

If you want to publish specific transient properties, you should list them within the `publishWithSnapshots()` call, as in the following example.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)



## JShell

```
opg4j> var prop1 = graph.createVertexProperty(PropertyType.INTEGER,
"prop1")
opg4j> prop.fill(0)
opg4j> var cost = graph.createEdgeProperty(PropertyType.DOUBLE, "cost")
opg4j> cost.fill(0d)
opg4j> graph.publishWithSnapshots(List.of(prop1), List.of(cost))
```

## Java

```
VertexProperty<Integer, Integer> prop1 =
graph.createVertexProperty(PropertyType.INTEGER, "prop1");
prop.fill(0);
EdgeProperty<Double> cost =
graph.createEdgeProperty(PropertyType.DOUBLE, "cost");
cost.fill(0d);
List<VertexProperty<Integer, Integer>> vertexProps =
Arrays.asList(prop);
List<EdgeProperty<Double>> edgeProps = Arrays.asList(cost);
graph.publishWithSnapshots(vertexProps, edgeProps);
```

## Python

```
VertexProperty<Integer, Integer> prop1 =
graph.createVertexProperty(PropertyType.INTEGER, "prop1")
prop.fill(0)
EdgeProperty<Double> cost =
graph.createEdgeProperty(PropertyType.DOUBLE, "cost")
cost.fill(0d)
List<VertexProperty<Integer, Integer>> vertexProps = Arrays.asList(prop)
List<EdgeProperty<Double>> edgeProps = Arrays.asList(cost)
graph.publishWithSnapshots(vertexProps, edgeProps)
```

 **Note:**

The published properties, like the original transient properties, are associated to the specific snapshot they had been created on, so **they are not visible on other snapshots**.

### Referencing a Published Graph from Another Session

Other sessions can reference a published graph by its name via the `getGraph()` method of the session object.

The following example references a published graph of `session1`, `myGraph`, in `session2`.

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var session2 = instance.createSession("session2")
opg4j> var graph2 = session2.getGraph(Namespace.PUBLIC, "myGraph")
```

## Java

```
PgxSession session2 = instance.createSession("session2");
PgxGraph graph2 = session2.getGraph(Namespace.PUBLIC, "myGraph");
```

## Python

```
session2 = pypgx.get_session("session2");
PgxGraph graph2 = session2.get_graph("myGraph")
```

---

`session2` can see only the published snapshot. If the graph has been published without snapshots, calls to the `getAvailableSnapshots()` method of `session2` return an empty queue.

Instead, if also the snapshots have been published, the call to `getGraph()` returns the most recent snapshot available. `session2` can see all the available snapshots via `getAvailableSnapshots()` and set a specific one via the `setSnapshot()` method of `PgxSession`.



### Note:

You must remember to release every graph you reference, when you do not need it anymore. See [Deleting a Graph](#) for more information.

### Publishing a Property

After publishing (a single snapshot or all of them), you can still publish transient properties individually:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> graph.getVertexProperty("prop1").publish()
opg4j> graph.getEdgeProperty("cost").publish()
```

## Java

```
graph.getVertexProperty("prop1").publish();
graph.getEdgeProperty("cost").publish();
```

## Python

```
graph.get_vertex_property("prop1").publish()
graph.get_edge_property("cost").publish()
```



### Note:

Published properties are associated to the specific snapshot they have been created on and thus visible only on that snapshot.

### Getting a Published Property in Another Session

Sessions referencing a published graph (with or without snapshots) can reference a published property via the usual `getVertexProperty` and `getEdgeProperty` calls of `PgxGraph`.

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var session2 = instance.createSession("session2")
opg4j> var graph2 = session2.getGraph(Namespace.PUBLIC, "myGraph")
opg4j> var vertexProperty = graph2.getVertexProperty("prop1")
opg4j> var edgeProperty = graph2.getEdgeProperty("cost")
```

## Java

```
PgxSession session2 = instance.createSession("session2");
PgxGraph graph2 = session2.getGraph(Namespace.PUBLIC, "myGraph");
VertexProperty<Integer, Integer> vertexProperty =
```

```
graph2.getVertexProperty("prop1");
EdgeProperty<Double> edgeProperty = graph2.getEdgeProperty("cost");
```

## Python

```
session2 = pygpx.get_session(session_name="session2")
graph2 = session2.get_graph("myGraph")
vertex_property = graph2.get_vertex_property("prop1")
edge_property = graph2.get_edge_property("cost")
```

---

`session2` now has a reference to the published graph of `session1` called `myGraph` and can reference its published properties via `myGraph` itself.

## 12.3 Publishing a Preloaded Graph

The publishing behavior for preloaded graphs can be controlled in the configuration. Unless a different behavior is configured, (only) the first loaded snapshot of a graph is published. Preloaded published graphs remain in memory even if they are not used by any session.

There are two options to control the publishing behavior:

- Set the optional flag `publish` to `true`, to publish only the graph but no future snapshots of the graph. This is the default behavior as the default value of this flag is `true`.
- Set the optional flag `publish_with_snapshots` to `true`, to publish the graph and all future snapshots of the graph. The default value is `false`.

Only one of these two flags can be set to `true` at a time. However, publishing the graph with snapshots does also publish the first version of the graph.

### Example 12-1 Sample Configuration File for Preloading Graphs

This example `pgx.conf` specifies two graphs for loading into memory during the graph server (PGX) startup-time. `my-graph` is published with snapshots while `my-other-graph` is published without snapshots.

```
{
 "enterprise_scheduler_config": {
 "analysis_task_config": {
 "max_threads": 32
 }
 },
 "preload_graphs": [
 {
 "path": "graph-configs/my-graph.bin.json",
 "name": "my-graph",
 "publish": false,
 "publish_with_snapshots": true
 },
 {
 "path": "graph-configs/my-other-graph.adj.json",
 "name": "my-other-graph" }
}
```

```
 }
],
 "authorization": [{
 "pgx_role": "GRAPH_DEVELOPER",
 "pgx_permissions": [{
 "preloaded_graph": "my-graph",
 "grant": "read"
 }],
 {
 "preloaded_graph": "my-other-graph",
 "grant": "read"
 }
]
},

]
}
```

The two preloaded graphs can be accessed as follows:

```
PgxGraph g1 = session.getGraph("my-graph"); //returns the most recent
available snapshot
PgxGraph g2 = session.getGraph("my-other-graph");
```

## 12.4 Deleting a Graph

In order to reduce the memory usage of the graph server (PGX), the session must drop the unused `PgxGraph` graph objects that it created via `PgxSession.getGraph()` by invoking the `destroy()` method.

This step not only destroys the specified graph, but all of its associated properties, including transient properties as well. In addition, all of the collections related to the graph instance (for example, a `VertexSet`) are also destroyed automatically. If a session holds multiple `PgxGraph` objects referencing the same graph, invoking `destroy()` on any of them will invalidate all the `PgxGraph` objects referencing that graph, making any operation on those objects fail.

For example:

- 
- [Java](#)
  - [Python](#)

### Java

```
PgxGraph graph1 = session.getGraph("myGraphName");

// graph2 references the same graph of graph1
PgxGraph graph2 = session.getGraph("myGraphName");

// Delete graph2
graph2.destroy();
```

```
// Both the following calls throw an exception, as both references are not
valid anymore
Set<VertexProperty<?, ?>> properties = graph1.getVertexProperties();
properties = graph2.getVertexProperties();
```

## Python

```
graph1 = session.get_graph("myGraphName")

graph2 references the same graph of graph1
graph2 = session.get_graph("myGraphName")

Delete graph2
graph2.destroy()

Both the following calls throw an exception, as both references are not
valid anymore
properties = graph1.get_vertex_properties()
properties = graph2.get_vertex_properties()
```

---

The same behavior occurs when multiple `PgxGraph` objects reference the same snapshot. Since a snapshot is effectively a graph, destroying a `PgxGraph` object referencing a certain snapshot invalidates all `PgxGraph` objects referencing the same snapshot, but does not invalidate those referencing other snapshots:

```
// Get a snapshot of "myGraphName"
PgxGraph graph1 = session.getGraph("myGraphName");

// graph2 and graph3 reference the same snapshot as graph1
PgxGraph graph2 = session.getGraph("myGraphName");
PgxGraph graph3 = session.getGraph("myGraphName");

// Assume another snapshot is created ...

// Make graph3 references the latest snapshot available
session.setSnapshot(graph3, PgxSession.LATEST_SNAPSHOT);
graph2.destroy();

// Both the following calls throw an exception, as both references are not
valid anymore
Set<VertexProperty<?, ?>> properties = graph1.getVertexProperties();
properties = graph2.getVertexProperties();

// graph3 is still valid, so the call succeeds
properties = graph3.getVertexProperties();
```



**Note:**

Even if a graph is destroyed by a session, the graph data may still remain in the server memory if the graph is currently shared by other sessions. In such a case, the graph may still be visible among the available graphs via `PgxSession.getGraphs()`.

As a safe alternative to manual destruction of each graph, the PGX API supports some implicit resource management features which allow developers to safely omit the `destroy()` call. See [Resource Management Considerations](#) for more information.

# 13

## Namespaces and Sharing

The in-memory graph server (PGX) supports separate namespaces that help you to organize your entities.

Each client session has its own session-private namespace and can choose any name without affecting other sessions. There is also a public namespace for published graphs (for example, published via the `publishWithSnapshots()` or the `publish()` methods).

Similarly, each published graph defines a public namespace for published properties as well as a private namespace per session. So different sessions can create properties with the same name on a published graph.

- [Defining Graph Names](#)
- [Retrieving Graphs by Name](#)
- [Checking Used Names](#)
- [Property Name Resolution and Graph Mutations](#)

### 13.1 Defining Graph Names

Graphs that are created in a session either through loading (for example, calling `readGraphWithProperties()`) or through mutations will take up a name in the session-private namespace. A graph will be placed in the public namespace only through publishing (that is, when calling the `publishWithSnapshots()` or the `publish()` methods). Publishing a graph will move its name from the session-private namespace to the public namespace.

There can only be one graph with a given name in a given namespace, but a name can be used in different namespaces to refer to different graphs. An operation that creates a new graph (for example, `readGraphWithProperties()`) will fail if the chosen name of the new graph already exists in the session-private namespace. Publishing a graph fails if there is already a graph in the public namespace with the same name.

### 13.2 Retrieving Graphs by Name

You can retrieve a graph by name by the following two ways:

- `getGraph(Namespace, String)`: with explicitly mentioning the namespace
- `getGraph(String)`: without explicitly mentioning the namespace

With `getGraph(Namespace, String)`, you need to provide the namespace (either session private or public). In this case, the graph will be looked up in the given namespace only.

With `getGraph(String)`, the provided name will be first looked up in the private namespace. If no graph with the given name is found there, then the graph name will be looked up in the public namespace. In other words, if a graph with the same name is defined in both the public and the private namespaces, `getGraph(String)` will return the private graph and you need to use `getGraph(Namespace, String)` to get hold of the public graph with that name.



## 13.3 Checking Used Names

To see the currently used names in a namespace you can use the [PgxSession.getGraphs\(Namespace\)](#) method, which will list all the names in the given namespace. The names in the returned collection can be used in a [getGraph\(Namespace, String\)](#) call to retrieve the corresponding `PgxGraph`.

## 13.4 Property Name Resolution and Graph Mutations

Property names behave in a similar way as graph names. All property names of a non-published graph are in the session-private namespace. Once a graph is published with [PgxGraph.publishWithSnapshots\(\)](#) or the [PgxGraph.publish\(\)](#) methods, its properties are published as well and their names move into the public namespace.

Once a graph is published, newly created properties will still be private to the session and their names will be in the private namespace. Those properties can be published individually with the [Property.publish\(\)](#) method, as long as no other property with the same name is already published for that graph.

Additionally, new private properties can be created with the same name of an already-published properties (since the names are part of separate namespaces). To handle such situations and retrieve the correct property, the PGX API offers the [getVertexProperty\(Namespace, String\)](#) and the [getEdgeProperty\(Namespace, String\)](#) methods, which allow specifying the namespace where the property name should be looked up.

Similar to graphs, if you search a property without specifying the namespace, the private namespace is searched first and if the property is not found, the search proceeds to the public namespace. This case applies for [getVertexProperty\(String\)](#) or the [getEdgeProperty\(String\)](#) methods and for PGQL queries.

Likewise, when a mutation on a graph reads or writes a property referred to by name and two properties exist with the same name, the property in the private namespace is selected. To override the default selection, some mutation mechanisms accept a collection of specific `Property` objects to be copied into the mutated graph. For example, such mechanism is supported for filter expressions. See [Creating Subgraphs](#) for more details.

# 14

## PGX Programming Guides

You can avail all the PGX functionalities through asynchronous Java APIs. Each asynchronous method has a synchronous equivalent, which blocks the caller thread until the server produces a response.

These APIs may perform one or any combination of:

- Complex, non-blocking Java applications on top of PGX
- Simple, sequential Java scripts executed by JShell
- ShellPerforming interactive graph analysis in the JShell

### Layers of PGX API

The PGX API is composed of a few different Java interfaces. Each interface provides a distinct layer of abstraction for PGX, as shown in the following table:

**Table 14-1 PGX API Interface**

Interface	Description
ServerInstance	The <code>ServerInstance</code> class encapsulates access to a PGX server instance and can be used to create sessions, start and stop the PGX engine, monitor the engine status and perform other administrative tasks. If the instance points to a remote instance, access to the administrative functions requires special authorization on the HTTP level by default.
PgxSession	A <code>PgxSession</code> represents an active user currently connected to an instance. Each session gets its own workspace on the server side which can be used to read graphs, create in-memory data structures, hold analysis results and custom algorithms. The <code>PgxSession</code> class provides various methods to create new transient data (currently collections). If a session is idling for too long, the PGX engine will automatically destroy it to ensure no resources are wasted.
PgxGraph	A <code>PgxGraph</code> represents a client-side handle to the graph data managed by the PGX server. A graph may contain an arbitrary amount of properties of type <code>VertexProperty</code> and/or <code>EdgeProperty</code> .

#### Note:

The PGX currently only supports non-partitioned graphs, meaning every vertex/edge has the same properties with the same names and types as all the other vertices/edges.

`PgxGraph` class provides various methods to create new transient data (including maps and collections) as well as graph mutation operations, such as undirecting, sorting and filtering.

**Table 14-1 (Cont.) PGX API Interface**

Interface	Description
<code>Analyst</code>	The <code>Analyst</code> API contains all of the built-in algorithms PGX provides. <code>Analyst</code> objects keep track of all the transient data they created during algorithm invocations to hold analysis results. Once an <code>Analyst</code> gets destroyed, all the results it created get freed on the server-side automatically.
<code>CompiledProgram</code>	The <code>CompiledProgram</code> class (PGX Algorithm API) encapsulates runtime-compiled custom algorithms and allows invocation of those algorithms using PGX data objects, such as <code>PgxGraph</code> or <code>VertexProperty</code> , as arguments.

Please see the [oracle.pgx.api](#) package in the Javadoc for more details.

- [Design of the Graph Server \(PGX\) API](#)  
This guide focuses on the design of the graph server (PGX) API.
- [Data Types and Collections in the Graph Server \(PGX\)](#)  
This guide provides you the list of the supported data types and collections in the graph server (PGX).
- [Handling Asynchronous Requests in Graph Server \(PGX\)](#)  
This guide explains in detail the asynchronous methods supported by the PGX API.
- [Graph Client Sessions](#)  
The graph server (PGX) assumes there may be multiple concurrent clients, and each client submits request to the shared PGX server independently.
- [Graph Mutation and Subgraphs](#)  
This guide discusses the several methods provided by the graph server (PGX) for mutating graph instances.
- [Graph Builder and Graph Change Set](#)
- [Managing Transient Data](#)  
This guide discusses how to handle transient properties and collections.
- [Graph Versioning](#)  
This guide describes the different ways to work with graph snapshots.
- [Labels and Properties](#)  
You can perform various actions on the graph property and label values by executing PGQL queries.
- [Filter Expressions](#)  
This guide explains the usage of filter expressions.
- [Advanced Task Scheduling Using Execution Environments](#)  
This guide shows how you can use the advanced scheduling features of the enterprise scheduler.
- [Admin API](#)  
This guide shows how to use the graph server (PGX) Admin API to inspect the server state including sessions, graphs, tasks, memory and thread pools.
- [PgxFrames Tabular Data-Structure](#)

## 14.1 Design of the Graph Server (PGX) API

This guide focuses on the design of the graph server (PGX) API.

The design of the PGX API reflects consideration of the following situations:

- Multiple clients may concurrently be accessing a single running instance of PGX, sharing its resources. Each client needs to maintain its own isolated workspace (session).
- Graph and property data can be large in size and therefore that data only resides on the server side.
- Some graph analysis may take a significant amount of time.
- Clients may not reside in the same address space (JVM) as PGX. Actually, clients may not even be Java applications.

### Client Sessions

In PGX, each client maintains its own session, an isolated, private workspace. Therefore, clients first have to obtain a `PgxSession` object from a `PGX ServerInstance` before they can perform any analysis.

### Asynchronous Execution

The PGX API is designed for asynchronous execution. That means that each computationally intensive method in the PGX API *immediately* returns a `PgxFuture` object without waiting for the request to finish. The `PgxFuture` class implements the `Future` interface, which can be used to retrieve the result of a computation at some point in the future.

#### Note:

The asynchronous execution aspect of this design facilitates multiple (remote) clients submitting requests to a single server. A request from one client may be queued up to wait until PGX resources become available. The asynchronous API allows the client (or calling thread) to work on other tasks until PGX completes the request.

### No Direct References

The PGX API does not return objects with direct reference to PGX internal objects (such as the graph or its properties) to the client. This is because:

- The client might not be in the same JVM as the server
- The graph instance might be shared by multiple clients

Instead, the PGX API only returns lightweight, stateless pointer objects to those objects. These pointer objects only holds the `ID(name)` of the server-side object to which they are pointing.

### Resource Management Considerations

The in-memory graph server (PGX), being an *in-memory* analytic engine, might allocate large amounts of memory to hold the graph data of clients. Therefore, it is important that client

sessions clean up their resources once they have ended. The PGX API supports several features to make this easier:

- Every object returned by the PGX API pointing to a server-side resource implements the `Destroyable` interface, which means all memory-consuming client-side objects can be destroyed the same way. For example:

```
PgxGraph myGraph = ...
myGraph.destroyAsync(); // request destruction of myGraph, don't
wait for response
try {
 myGraph.destroy(); // blocks caller thread until destruction
was done
} catch (ExecutionException e) {
 // destruction failed
}
```

- `Destroyable` extends [AutoClosable](#), so users can leverage Java's built-in resource management syntax:

```
try (PgxGraph myGraph = session.readGraphWithProperties(config)) {
 // do something with myGraph
}
// myGraph is destroyed
```

- **Session time out.** In some cases, the PGX server will remove the session and all its data automatically. This can occur when a client fails to destroy either the data or its session, or if it does not hear from the session after a configurable timeout. See [Configuration Parameters for the Graph Server \(PGX\) Engine](#) for more information to configure timeout parameters.

## 14.2 Data Types and Collections in the Graph Server (PGX)

This guide provides you the list of the supported data types and collections in the graph server (PGX).

### Primitive Data Types

The following section explains the primitive data types supported by the graph server (PGX) and their limitations.

PGX supports the following primitive data types.:

- **Numeric Types:** `integer`, `long`, `float`, and `double`. These types have the same size, range and precision of the corresponding Java primitive data type.
- **Boolean Type:** The `boolean` data type has only two possible values, `true` and `false`. As with Java and C++, its size is not precisely defined.
- **String:** `String` is a primitive data type in PGX. PGX follows the Java conventions for `String` representation.
- **Datetime Types:** `date`, `time`, `timestamp`, `time with time zone`, and `timestamp with time zone`. These types correspond to the Java types shown in [Table 14-2](#) from the standard library package `java.util.time`.

- **Vertex and Edge:** The type `vertex` or `edge` of the graph itself is a proper type in PGX.

 **Note:**

- `vertex` and `edge` is itself a valid primitive data type. For instance, in a path-finding algorithm, each `vertex` can have a temporary property `predecessor` that stores which incoming neighbor is the predecessor vertex in the path. Such a property would have the type `vertex`.
- `local_date` must be used instead of `date` in the graph configuration file. See [Using Datetime Data Types](#) for more examples on usage of datetime data types.

All properties and scalar variables must be one of the above preceding data types. See [Managing Transient Data](#) for more information on handling transient properties and scalar variables.

The following table presents the overview of the supported data types, their integration in different languages and APIs and their minimum and maximum value limitations.

 **Note:**

- For float and double types, the smallest absolute value is included in the table, the minimum value is the negative of maximum value for these types.
- For string values, PGX supports arbitrary long strings.

**Table 14-2 Overview of Data types**

Data Type	Loading & Storing	PGX Java API	PGQL and Filter Expression	Minimum Value Limitation	Maximum Value Limitation
string	string	String	STRING	-	-
int/integer	int/integer	int	INT/INTEGER	-2147483648	2147483647
long	long	long	LONG	-9223372036854775808	9223372036854775807
float	float	float	FLOAT	1.4E-45	3.4028235e+38
double	double	double	DOUBLE	4.9E-324	1.7976931348623157E308
boolean	boolean	boolean	BOOLEAN	-	-
date	local_date	LocalDate	DATE	-5877641-06-23	5881580-07-11
time	time	LocalTime	TIME	00:00:00.000	23:59:59.999
timestamp	timestamp	LocalDateTime	TIMESTAMP	-292275055-05-17 00:00:00.000	292278994-08-17 07:12:55.807
time with time zone	time_with_timezone	OffsetTime	TIME WITH TIME ZONE	00:00:00.000+18:00	23:59:59.999-18:00

**Table 14-2 (Cont.) Overview of Data types**

Data Type	Loading & Storing	PGX Java API	PGQL and Filter Expression	Minimum Value Limitation	Maximum Value Limitation
timestamp with time zone	timestamp_with_timezone	OffsetDateTime	TIMESTAMP WITH TIME ZONE	-292275055-05-17 00:00:00.000+18:00	292278994-08-17 07:12:55.807-18:00
vertex	-	PgxVertex	-	-	-
edge	-	PgxEdge	-	-	-

### Collections

The in-memory graph server (PGX) supports three different collection types: `sequence`, `set` and `order`. All of these collections can contain values of the `vertex` type, but each has different semantics regarding uniqueness and preserving the order of its elements:

- **Sequence:** a `sequence` works basically like a list. It preserves the order of the elements added to it, and the same element can appear multiple times.
- **Set:** a `set` can contain the same value once at the most. Adding a value that is already in the `set` will have no effect. `set` does not preserve the order of the elements it contains.
- **Order:** just like the `set`, the `order` collection will contain each element once at the most. But the `order` preserves the order of the elements inserted into it (that is, it is a FIFO data structure).

See [Collection Data Types](#) for examples on creation and usage of the different collections.

### Immutable Collections

Some operations, like `PgxGraph.getVertices()` and `PgxGraph.getEdges()` return immutable collections. These collections behave like normal collections, but cannot be modified by operations like `addAll` or `removeAll` and `clear`.

An immutable collection can be transformed into a mutable collection by using the `toMutable` method, which returns a mutable copy of the collection. If `toMutable` is called on a collection that is already mutable, the method has the same result as the method `clone`.

To check if a collection is mutable, use the `isMutable` method.

### Maps

PGX provides the following two kinds of maps:

- Graph-bound maps can hold mappings between types in `PropertyType`. This is the kind of maps to use if the key or value types are graph-related like `VERTEX` and `EDGE` otherwise using session-bound maps is recommended.
- Session-bound maps can map between non graph-related types and are directly bound to the session.

See [Map Data Types](#) for examples on creation and usage of maps.

- [Using Collections and Maps](#)
- [Using Datetime Data Types](#)

## 14.2.1 Using Collections and Maps

This section explains with examples, the creation and usages of collections and maps.

You must first create a session before getting started with the collection and map data types.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
cd /opt/oracle/graph/
./bin/opg-jshell // starting the shell will create an implicit session
```

### Java

```
import oracle.pgx.api.*;
...
PgxSession session=Pgx.createSession("<session_name>");
```

### Python

```
from pypgx import get_session
session = get_session(session_name="<session_name>")
```

- 
- [Collection Data Types](#)
  - [Map Data Types](#)

### 14.2.1.1 Collection Data Types

The in-memory graph server (PGX) defines two types of collections:

- **Graph-bound collections:** such as vertex and edge collections. These collections belong to the graph.
- **Session-bound collections:** belong to the session.
- [Graph-Bound Collections](#)
- [Session-Bound Collections](#)



### 14.2.1.1.1 Graph-Bound Collections

The following describes the usage of graph-bound collections.

You must first load the graph to work with vertex and edge collections as shown in [Loading a Graph Into the Graph Server \(PGX\)](#).

#### Vertex Collections

You can create a vertex collection as shown in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
v0 = graph.getVertex(100) // 'graph' is the loaded graph object. '100'
-> '103' are vertex ids that supposedly
v1 = graph.getVertex(101) // exist in the graph
v2 = graph.getVertex(102)
v3 = graph.getVertex(103)

myVertexSet = graph.createVertexSet("myVertexSet") // A name is
automatically generated if none given
myVertexSet.add(v0) // Adds vertex
'v0' to the set
myVertexSet.addAll([v1, v2, v3]) // Supports
variadic parameter as well: myVertexSet.addAll(v1, v2, v3)
```

#### Java

```
import java.util.Arrays;
import oracle.pgx.api.*;
...
PgxVertex v0 = graph.getVertex(100);
PgxVertex v1 = graph.getVertex(101);
PgxVertex v2 = graph.getVertex(102);
PgxVertex v3 = graph.getVertex(103);

VertexSet myVertexSet = graph.createVertexSet("myVertexSet"); // A
name is automatically generated if none given
myVertexSet.add(v0);
myVertexSet.addAll(Arrays.asList(v1, v2, v3));
```

## Python

```
...
v0 = graph.get_vertex(100)
v1 = graph.get_vertex(101)
v2 = graph.get_vertex(102)
v3 = graph.get_vertex(103)

my_vertex_set = graph.create_vertex_set("myVertexSet")
my_vertex_set.add(v0)
my_vertex_set.add_all([v1,v2,v3])
```

---

### Edge Collections

You can create an edge collection as shown in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
e0 = graph.getEdge(100) // 'graph' is the loaded graph object. '100' ->
'103' are edge ids that supposedly
e1 = graph.getEdge(101) // exist in the graph
e2 = graph.getEdge(102)
e3 = graph.getEdge(103)

myEdgeSequence = graph.createEdgeSequence("myEdgeSequence")
myEdgeSequence.add(e0)
myEdgeSequence.addAll([e1, e2, e3])
```

### Java

```
import java.util.Arrays;
import oracle.pgx.api.*;
...
PgxEdge e0 = graph.getEdge(100);
PgxEdge e1 = graph.getEdge(101);
PgxEdge e2 = graph.getEdge(102);
PgxEdge e3 = graph.getEdge(103);

EdgeSequence myEdgeSequence = graph.createEdgeSequence("myEdgeSequence");
myEdgeSequence.add(e0);
myEdgeSequence.addAll(Arrays.asList(e1, e2, e3));
```

## Python

```
e0 = graph.get_edge(100)
e1 = graph.get_edge(101)
e2 = graph.get_edge(102)
e3 = graph.get_edge(103)

my_edge_sequence = graph.create_edge_sequence("my_edge_sequence")
my_edge_sequence.add(e0)
my_edge_sequence.add_all([e1, e2, e3])
```

---

### 14.2.1.1.2 Session-Bound Collections

You can create and manipulate collections directly in the session without the need for a graph. Session-bound collections can be further passed as parameters to graph algorithms or used like any other collection object. The following sub-sections describe the currently supported types for these collections.

#### Scalar Collections

Scalar collections contain simple data types like `Integer`, `Long`, `Float`, `Double` and `Boolean`. They can be managed by the `PgxSession` APIs:

#### Creation of a Scalar Collection

You can use `createSet()` and `createSequence()` methods to create a scalar collection as shown in the following code:

- 
- [JShell](#)
  - [Java](#)

### JShell

```
myIntSet = session.createSet(PropertyType.INTEGER, "myIntSet")
myDoubleSequence = session.createSequence(PropertyType.DOUBLE) // A
name will be automatically generated if none is provided.
println myDoubleSequence.getName() //
Display the generated name.
```

### Java

```
import oracle.pgx.api.*;
import oracle.pgx.common.types.*;
...
ScalarSet myIntSet = session.createSet(PropertyType.INTEGER,
"myIntSet");
ScalarSequence myDoubleSequence =
```

```
session.createSequence(PropertyType.DOUBLE);
System.out.println(myDoubleSequence.getName());
```

---

## Run Operations on a Scalar Collection

You can run several operations on a scalar collection as shown in the following code:

- [JShell](#)
- [Java](#)

### JShell

```
myIntSet.add(10)
myIntSet.addAll([0, 1, 2, 3, 4, 5, 6, 7, 8, 9])
myIntSet.addAll([0,1,2]) // Element uniqueness. This
operation has no effect on the set.
println myIntSet

myIntSet.contains(1) // Checks the presence of an element.
This code returns `true`.
myIntSet.remove(10)
myIntSet.removeAll([4, 5, 6, 7, 8, 9]) // Leaves only elements `0, 1, 2, 3`.
println myIntSet
```

### Java

```
import java.util.Arrays;
import oracle.pgx.api.*;
...
myIntSet.add(10);
myIntSet.addAll(Arrays.asList(0, 1, 2, 3, 4, 5, 6, 7, 8, 9));
myIntSet.addAll(Arrays.asList(0, 1, 2));

myIntSet.contains(1); // Returns `true`.
myIntSet.remove(10);
myIntSet.removeAll(Arrays.asList(4, 5, 6, 7, 8, 9));
```

---

## Traversal of a Scalar Collection

You can traverse a scalar collection either using an iterator or using the new [Stream](#) API. You can add elements of a sequence to a set, traverse a sequence and filter out elements not required, and then add the rest to another scalar collection.

---

- [JShell](#)
- [Java](#)

## JShell

```
myIntSet.forEach({x -> print x + "\n"})
myIntSet.stream().filter({x -> x % 2 == 0}).forEach({x ->
myDoubleSequence.add(x)})
println myDoubleSequence
```

## Java

```
import java.util.Iterator;
import java.util.stream.Stream;
import oracle.pgx.api.*;
...
myIntSet.forEach(x -> System.out.println(x));
myIntSet.stream().filter(x -> x % 2 ==
0).forEach(myDoubleSequence::add);
```

---

### 14.2.1.2 Map Data Types

The in-memory graph server (PGX) defines two types of maps:

- **Graph-bound maps:** These maps support any key or value type and are created using a graph object.
- **Session-bound maps:** Keys or values in these maps are of any type except from graph-related types (that is, vertices or edges). These maps belong to the session.
- [Graph-Bound Maps](#)
- [Session-Bound Maps](#)

#### 14.2.1.2.1 Graph-Bound Maps

Some data types like `VERTEX` or `EDGE` depend on the graph. Consequently, mappings involving these data types also depend on the graph. PGX provides `PgxGraph` and `PgxMap` APIs to manage such maps.

The following describes the usage of graph-bound maps.

You must first load the graph to work with vertex and edge maps.

You can create a graph-bound map using vertices as keys as shown in the following code:

- 
- [JShell](#)
  - [Java](#)

- Python

## JShell

```
v0 = graph.getVertex(100)
v1 = graph.getVertex(101)
v2 = graph.getVertex(102)
v3 = graph.getVertex(103)

vertexToLongMap = graph.createMap(PropertyType.VERTEX, PropertyType.LONG,
"vertexToLongMap")
vertexToLongMap.put(v0, v0.getDegreeAsync().get())
vertexToLongMap.put(v1, v1.getDegreeAsync().get())
vertexToLongMap.put(v2, v2.getDegreeAsync().get())
vertexToLongMap.put(v3, v3.getDegreeAsync().get())
```

## Java

```
import java.util.Arrays;
import oracle.pgx.api.*;
...
PgxVertex v0 = graph.getVertex(100);
PgxVertex v1 = graph.getVertex(101);
PgxVertex v2 = graph.getVertex(102);
PgxVertex v3 = graph.getVertex(103);

PgxMap<PgxVertex, Long> vertexToLongMap =
graph.createMap(PropertyType.VERTEX, PropertyType.LONG, "vertexToLongMap");
vertexToLongMap.put(v0, v0.getDegree());
vertexToLongMap.put(v1, v1.getDegree());
vertexToLongMap.put(v2, v2.getDegree());
vertexToLongMap.put(v3, v3.getDegree());
```

## Python

```
v0 = graph.get_vertex(100)
v1 = graph.get_vertex(101)
v2 = graph.get_vertex(102)
v3 = graph.get_vertex(103)

vertex_to_long_map = graph.create_map("vertex", "long", "vertex_to_long_map")
vertex_to_long_map.put(v0, v0.degree)
vertex_to_long_map.put(v1, v1.degree)
vertex_to_long_map.put(v2, v2.degree)
vertex_to_long_map.put(v3, v3.degree)
```

---

You can create graph-bound maps using edges as keys as shown in the following code:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
e0 = graph.getEdge(100)
e1 = graph.getEdge(101)
e2 = graph.getEdge(102)
e3 = graph.getEdge(103)

edgeToVertexMap = graph.createMap(PropertyType.EDGE,
PropertyType.VERTEX, "edgeToVertexMap")
edgeToVertexMap.put(e0, e0.getSource())
edgeToVertexMap.put(e1, e1.getSource())
edgeToVertexMap.put(e2, e2.getSource())
edgeToVertexMap.put(e3, e3.getSource())
```

## Java

```
import java.util.Arrays;
import oracle.pgx.api.*;
...
PgxE edge e0 = graph.getEdge(100);
PgxE edge e1 = graph.getEdge(101);
PgxE edge e2 = graph.getEdge(102);
PgxE edge e3 = graph.getEdge(103);

PgxMap<PgxE, PgxVertex> edgeToVertexMap =
graph.createMap(PropertyType.EDGE, PropertyType.VERTEX,
"edgeToVertexMap");
edgeToVertexMap.put(e0, e0.getSource());
edgeToVertexMap.put(e1, e1.getSource());
edgeToVertexMap.put(e2, e2.getSource());
edgeToVertexMap.put(e3, e3.getSource());
```

## Python

```
e0 = graph.get_edge(100)
e1 = graph.get_edge(101)
e2 = graph.get_edge(102)
e3 = graph.get_edge(103)

edge_to_long_map = graph.create_map("edge", "long",
"edge_to_long_map")
edge_to_long_map.put(e0, e0.source)
edge_to_long_map.put(e1, e1.source)
edge_to_long_map.put(e2, e2.source)
edge_to_long_map.put(e3, e3.source)
```

**Note:**

If you destroy the graph you will lose the map. Consider using a session-bound maps instead if your map does not involve any graph-related key or value type.

### 14.2.1.2.2 Session-Bound Maps

You can directly create maps in the session. But, you cannot use any graph-related data type as the map key or value type. Session-bound maps can be further passed as parameters to graph algorithms or used like any other map object. They are managed by `PgxSession` and `PgxMaps` APIs.

Scalar collections contain simple data types like `Integer`, `Long`, `Float`, `Double` and `Boolean`. They can be managed by the `PgxSession` APIs.

#### Creation of a Session-bound Map

You can use `createMap()` method and its overloads to create a session-bound map.

- [JShell](#)
- [Java](#)

#### JShell

```
intToDouble = session.createMap(PropertyType.INTEGER, PropertyType.DOUBLE,
"intToDouble")
intToTime = session.createMap(PropertyType.INTEGER, PropertyType.TIME) // A
name will be automatically generated.
println intToTime.getName()
println intToTime.getSessionId()
println intToTime.getGraph() //
`null`: Not bound to a graph.
println intToTime.getKeyType()
println intToTime.getValueType()
```

#### Java

```
import java.time.LocalDateTime;
import oracle.pgx.api.*;
import oracle.pgx.common.types.*;
...
PgxMap<Integer, Double> intToDouble =
session.createMap(PropertyType.INTEGER, PropertyType.DOUBLE, "intToDouble");
PgxMap<Integer, LocalDateTime> intToTime =
session.createSequence(PropertyType.INTEGER, PropertyType.TIME);
System.out.println(intToTime.getName());
System.out.println(intToTime.getSessionId());
```



```
System.out.println(intToTime.getGraph()); // `null`: Not bound to a
graph.
System.out.println(intToTime.getKeyType());
System.out.println(intToTime.getValueType());
```

---

## Run Operations on a Session-bound Map

You can run important operations such as setting, removing and checking existence of entries on a session-bound map as shown in the following code:

- 
- [JShell](#)
  - [Java](#)

### JShell

```
intToDouble.put(0, 0.314)
intToDouble.put(1, 3.14)
intToDouble.put(2, 31.4)
intToDouble.put(3, 314)

println intToDouble.size() // 4
println intToDouble.get(1)
println intToDouble.get(3)
println intToDouble.get(10) // null

println intToDouble.containsKey(0) // `true`
intToDouble.remove(0)
println intToDouble.containsKey(0) // `false`
println intToDouble.containsKey(10) // `false`
intToDouble.remove(10)
println intToDouble.containsKey(10) // `false`

println intToDouble.put(1, 999) // previous mapped value
(`3.14`) is replaced by `999`
intToDouble.destroy()
```

### Java

```
import java.util.Arrays;
import oracle.pgx.api.*;

...

intToDouble.put(0, 0.314);
intToDouble.put(1, 3.14);
intToDouble.put(2, 31.4);
intToDouble.put(3, 314);
```

```
System.out.println(intToDouble.size()); // 4
System.out.println(intToDouble.get(1));
System.out.println(intToDouble.get(3));
System.out.println(intToDouble.get(10)); // null

System.out.println(intToDouble.containsKey(0)); // `true`
intToDouble.remove(0);
System.out.println(intToDouble.containsKey(0)); // `false`
System.out.println(intToDouble.containsKey(10)); // `false`
intToDouble.remove(10);
System.out.println(intToDouble.containsKey(10)); // `false`

System.out.println(intToDouble.put(1, 999)); // previous mapped value
(`3.14`) is replaced by `999`
intToDouble.destroy();
```

---

### Traversal of a Session-bound Map

You can traverse a session-bound map, using `entries()` method to get an iterable of map entries and `keys()` method to get an iterable of map keys.

- 
- [JShell](#)
  - [Java](#)

### JShell

```
intToDouble.entries().forEach {it -> println (it)}
intToDouble.keys().forEach {it -> println (it)}
```

### Java

```
import java.util.Iterable;
import java.util.stream.Stream;
import oracle.pgx.api.*;
...
Iterable<Map.Entry> entries = intToDouble.entries();
entries.forEach(System.out::println);
Iterable<Map.Entry> keys = intToDouble.keys();
keys.forEach(System.out::println);
```

---

## 14.2.2 Using Datetime Data Types

This section explains in detail working of datetime data types such as date, time and timestamp.

## Overview of Datetime Data Types in In-Memory Graph Server (PGX)

Table 14-3 presents the overview of the five `datetime` data types supported by PGX along with example values.



### Note:

PGX also supports custom format specification when loading data into PGX.

**Table 14-3 Overview of Datetime Data Types in PGX**

Data Type	Loading and Storing	PGX Java API	PGQL and Filter Expression	Example Value-1	Example Value-1
<code>date</code>	<code>local_date</code>	<code>LocalDate</code>	<code>DATE</code>	2001-01-29	2018-10-08
<code>time</code>	<code>time</code>	<code>LocalTime</code>	<code>TIME</code>	10:15	10:30:01.000
<code>timestamp</code>	<code>timestamp</code>	<code>LocalDateTime</code>	<code>TIMESTAMP</code>	2001-01-29 10:15	2018-10-08 10:30:01.000
<code>time with time zone</code>	<code>time_with_timezone</code>	<code>OffsetTime</code>	<code>TIME WITH TIME ZONE</code>	10:15+01:00	10:30:01.000-08:00
<code>timestamp with time zone</code>	<code>timestamp_with_timezone</code>	<code>OffsetDateTime</code>	<code>TIMESTAMP WITH TIME ZONE</code>	2001-01-29 10:15+01:00	2018-10-08 10:30:01.000-08:00

- [Loading Datetime Data](#)
- [Specifying Custom Datetime Formats](#)
- [APIs for Accessing Datetime Data](#)
- [Querying Datetime Data Using PGQL](#)
- [Accessing Datetimes from PGQL Result Sets](#)

### 14.2.2.1 Loading Datetime Data

You must first load a graph to work with datetime data. See [Loading a Graph Into the Graph Server \(PGX\)](#) for more information on graph loading.

The following example shows how to load a graph that has three vertices representing persons and zero edges.

#### Example 14-1 Loading Datetime Data

1. Create an `EDGE_LIST` file `persons.edge_list` as shown:

```
1*Judy,1989-01-15,1989-01-15 10:15-08:00
2*Klara,2001-01-29,2001-01-29 21:30-08:00
3*Pete,1995-08-01,1995-08-01 03:00-08:00
```

2. Create a corresponding graph configuration file `persons.edge_list.json` as shown:

```
{
 "format": "edge_list",
 "uri": "persons.edge_list",
 "vertex_id_type": "long",
 "vertex_props": [
 {
 "name": "name",
 "type": "string"
 },
 {
 "name": "date_of_birth",
 "type": "local_date"
 },
 {
 "name": "timestamp_of_birth",
 "type": "timestamp_with_timezone",
 "format": ["yyyy-MM-dd H[H]:m[m]:s[s]][XXX]"
 }
],
 "edge_props": [
],
 "separator": ","
}
```

3. You can now load the data as shown in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var graph = session.readGraphWithProperties("persons.edge_list.json",
"people_graph")
```

## Java

```
import oracle.pgx.api.*;
...
PgxGraph graph =
session.readGraphWithProperties("persons.edge_list.json", "people_graph");
```

## Python

```
graph =
session.read_graph_with_properties("persons.edge_list.json", graph_name=
"people_graph")
```

---

### 14.2.2.2 Specifying Custom Datetime Formats

You can also manually specify the datetime format(s) of your data.

By default, PGX tries to parse datetime values using a set of predefined formats. If this fails, an exception like the following is thrown:

```
property timestamp_of_birth: could not parse value at line 1 for
property of temporal type OffsetDateTime using any of the given formats
```

In such a case, you can custom format the datetime data.

There are two ways of specifying datetime formats:

- on a *per-property* basis
- on a *per-type* basis

#### Property-Specific Datetime format:

You can custom format the property `timestamp_of_birth` used in [Example 14-1](#) to the format `yyyy-MM-dd H[H]:m[m][:s[s]][XXX]` as shown:

#### Example 14-2 Specifying Property-Specific Datetime format:

```
{
 "name": "timestamp_of_birth",
 "type": "timestamp_with_timezone",
 "format": ["yyyy-MM-dd H[H]:m[m][:s[s]][XXX]"]
}
```

where `yyyy-MM-dd H[H]:m[m][:s[s]][XXX]` specifies that the timestamp values consist of:

- a four-digit year
- a hyphen followed by a two-digit month
- a hyphen followed by a two-digit day
- a space
- an hour, specified as either one or two digits
- a colon followed by a minute, specified as either one or two digits
- an *optional* part that consists of a colon followed by a second that is specified as either one or two digits
- an *optional* timezone

 **Note:**

- `H[H]:m[m]` allows the value `01:15` as well as the value `1:15`.
- `yyyy-MM-dd` allows the value `1989-01-15` but not the value `1989-1-15`. However, if two-digit months and days are needed, a format like `yyyy-M[M]-d[d]` can be used.

Also the format specification takes a *list* of formats. In the preceding example, the list contains only a single format, but you may specify any number of formats. If more than one format is specified, then when parsing the datetime data, the formats are tried from left to right until parsing succeeds. In this way, you can even load data that contains a mixture of values in different formats.

**Type-Specific Datetime format:**

You can also specify datetime formats on a *per-type* basis. This is useful in cases when there are multiple properties that have the same type as well as the same format because you will only need to specify the datetime format only once.

In case of the per-type specification, the format is used for each vertex or edge property that has the particular type.

The following example shows two type-specific formats (`local_date_format` and `timestamp_with_timezone_format`):

**Example 14-3 Specifying Type-Specific Datetime format:**

```
...
 "edge_props": [
],
 "separator": ",",
 "local_date_format": ["yyyy-MM-dd"],
 "timestamp_with_timezone_format": ["yyyy-MM-dd H[H]:m[m][:s[s]][XXX]"]
}
```

In the example, properties of type `date` (`local_date`) have the format `yyyy-MM-dd` while properties of type `timestamp with time zone` (`timestamp_with_timezone`) have the format `yyyy-MM-dd H[H]:m[m][:s[s]][XXX]`.

 **Note:**

Property-specific formats always overrides type-specific formats. If you specify a type-specific format, and the property of the particular type also has a property-specific format, then only the property-specific format is used to parse the datetime data.

### 14.2.2.3 APIs for Accessing Datetime Data

The in-memory graph server (PGX) uses the new [Java 8 temporal data types](#) for accessing datetime data through the Java API:

- `date` in PGX maps to `LocalDate` in Java
- `time` in PGX maps to `LocalTime` in Java
- `timestamp` in PGX maps to `LocalDateTime` in Java
- `time with time zone` in PGX maps to `OffsetTime` in Java
- `timestamp with time zone` in PGX maps to `OffsetDateTime` in Java

You can retrieve a date as shown in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var dateOfBirthProperty =
graph.getVertexProperty("date_of_birth")
opg4j> var birthdayOfJudy = dateOfBirthProperty.get(1)
```

## Java

```
import java.time.LocalDate;
import oracle.pgx.api.*;
...
VertexProperty<LocalDate> dateOfBirthProperty =
graph.getVertexProperty("date_of_birth");
LocalDate birthdayOfJudy = dateOfBirthProperty.get(1);
```

## Python

```
date_of_birth_property = graph.get_vertex_property("date_of_birth")
birthday_of_judy = date_of_birth_property.get(1)
```

---

### 14.2.2.4 Querying Datetime Data Using PGQL

You can perform various operations such as *extracting* values from datetimes, *comparing* datetime values, and, *converting* between different datetime types. on datetime data using PGQL.

The following are example PGQL queries that show different operations that involve datetime data:

## Retrieving Datetime Properties

The following query retrieves the `date_of_birth` and `timestamp_of_birth` properties from the all the persons in the graph.

```

SELECT n.name AS name, n.date_of_birth AS birthday, n.timestamp_of_birth
AS timestamp
FROM MATCH (n) ON people_graph
ORDER BY birthday

```

The result of the query is as follows:

```

+-----+
| name | birthday | timestamp |
+-----+
| Judy | 1989-01-15 | 1989-01-15T10:15-08:00 |
| Pete | 1995-08-01 | 1995-08-01T03:00-08:00 |
| Klara | 2001-01-29 | 2001-01-29T21:30-08:00 |
+-----+

```

## Comparing Datetime Values

The following query provides an overview of persons who are older than other persons in the graph:

```

SELECT n.name AS person1, 'is older than' AS relation, m.name AS person2
FROM MATCH (n) ON people_graph, (m) ON people_graph
WHERE n.date_of_birth > m.date_of_birth
ORDER BY person1, person2

```

The result of the query is as follows:

```

+-----+
| person1 | relation | person2 |
+-----+
| Klara | is older than | Judy |
| Klara | is older than | Pete |
| Pete | is older than | Judy |
+-----+

```

## Extracting Values from Datetimes

The following query extracts the year, month, and day from the `date_of_birth` values:

```

SELECT n.name AS name
, n.date_of_birth AS dob
, EXTRACT(YEAR FROM n.date_of_birth) AS year
, EXTRACT(MONTH FROM n.date_of_birth) AS month
, EXTRACT(DAY FROM n.date_of_birth) AS day
FROM MATCH (n) ON people_graph
ORDER BY name

```



The result of the query is as follows:

```
+-----+
| name | dob | year | month | day |
+-----+
| Judy | 1989-01-15 | 1989 | 1 | 15 |
| Klara | 2001-01-29 | 2001 | 1 | 29 |
| Pete | 1995-08-01 | 1995 | 8 | 1 |
+-----+
```

### Converting Between Different Types of Datetime Values

The following query converts the `timestamp_of_birth` property into values of the following three datetime types:

- a timestamp (without time zone)
- a time with time zone
- a time (without time zone)

```
SELECT n.name AS name
 , n.timestamp_of_birth AS original_timestamp
 , CAST(n.timestamp_of_birth AS TIMESTAMP) AS utc_timestamp
 , CAST(n.timestamp_of_birth AS TIME WITH TIME ZONE) AS
timezoned_time
 , CAST(n.timestamp_of_birth AS TIME) AS utc_time
FROM MATCH (n) ON people_graph
ORDER BY original_timestamp
```

The result of the query is as follows:

```
+-----+
-----+
| name | original_timestamp | utc_timestamp | timezoned_time |
| utc_time |
+-----+
-----+
| Judy | 1989-01-15T10:15-08:00 | 1989-01-15T18:15 | 10:15-08:00 |
| 18:15 |
| Pete | 1995-08-01T03:00-08:00 | 1995-08-01T11:00 | 03:00-08:00 |
| 11:00 |
| Klara | 2001-01-29T21:30-08:00 | 2001-01-30T05:30 | 21:30-08:00 |
| 05:30 |
+-----+
-----+
```

#### 14.2.2.5 Accessing Datetimes from PGQL Result Sets

You can use the following APIs for retrieving datetime values from PGQL result sets.

```
LocalDate getDate(int elementIdx)
LocalDate getDate(String variableName)
LocalTime getTime(int elementIdx)
```

```
LocalTime getTime(String variableName)
LocalDateTime getTimestamp(int elementIdx)
LocalDateTime getTimestamp(String variableName)
OffsetTime getTimeWithTimezone(int elementIdx)
OffsetTime getTimeWithTimezone(String variableName)
OffsetDateTime getTimestampWithTimezone(int elementIdx)
OffsetDateTime getTimestampWithTimezone(String variableName)
```

The following example prints the birthdays of all the persons in the graph is as follows:

- 
- [JShell](#)
  - [Java](#)

## JShell

```
opg4j> var resultSet = session.queryPgql("""
 SELECT n.name, n.date_of_birth
 FROM MATCH (n) ON people_graph
 ORDER BY n.name
 """)
opg4j> while (resultSet.next()) {
...> System.out.println(resultSet.getString(1) + " has birthday " +
resultSet.getDate(2));
...> }
opg4j> resultSet.close()
```

## Java

```
import java.time.LocalDate;
import oracle.pgx.api.*;
...
PgqlResultSet resultSet = session.queryPgql(
 " SELECT n.name, n.date_of_birth\n" +
 " FROM MATCH (n) ON people_graph\n" +
 " ORDER BY n.name");

while (resultSet.next()) {
 System.out.println(resultSet.getString(1) + " has birthday " +
resultSet.getDate(2));
}

resultSet.close();
```

---

The result of the query is as follows:

```
Judy has birthday 1989-01-15
Klara has birthday 2001-01-29
Pete has birthday 1995-08-01
```

In addition to the Java types from the new `java.time` package, the legacy `java.util.Date` is also supported through the following APIs:

```
Date getLegacyDate(int elementIdx)
Date getLegacyDate(String variableName)
```

 **Note:**

The legacy `java.util.Date` can store dates, times, as well as timestamps, so these two APIs can be used for accessing values of any of the five datetime types.

## 14.3 Handling Asynchronous Requests in Graph Server (PGX)

This guide explains in detail the asynchronous methods supported by the PGX API.

The PGX API is designed to be asynchronous. This means that all of its core methods ending with `Async` **do not** block the caller thread until the request is completed. Instead, a `PgxFuture` object is instantly returned.

You can perform the following three actions on the returned `PgxFuture` object:

- Block
- Chain
- Cancel
- [Blocking Operation](#)
- [Chaining Operation](#)
- [Cancelling Operation](#)
- [Handling Concurrent Asynchronous Operations](#)

### 14.3.1 Blocking Operation

You can easily get the result by calling the `get()` method on the `PgxFuture`. The `get()` blocks the caller thread until the result is available:

```
PgxFuture<PgxSession> sessionPromise = instance.createSessionAsync("my-
session");
try {
 // block caller thread
 PgxSession session = sessionPromise.get();
}
```

```

 // do something with session
 ...
 } catch (InterruptedException e) {
 // caller thread was interrupted while waiting for result
 } catch (ExecutionException e) {
 // an exception was thrown during asynchronous computation
 Throwable cause = e.getCause(); // the actual exception is nested
 }
}

```

PGX provides blocking convenience methods for every *Async* method, which calls the `get()` method. Typically, those methods have the same name as the asynchronous method they wrap, but without the *Async* suffix. For example, the preceding code snippet is equal to:

```

try {
 // block caller thread
 PgxFuture session = instance.createSession("my-session");
 // do something with session
 ...
} catch (InterruptedException e) {
 // caller thread was interrupted while waiting for result
} catch (ExecutionException e) {
 // an exception was thrown during asynchronous computation
 Throwable cause = e.getCause(); // the actual exception is nested
}

```

## 14.3.2 Chaining Operation

The in-memory graph server (PGX) ships a version of Java 8's `CompletableFuture` named `PgxFuture`, a monadic enhancement of the `Future` interface.

The `CompletableFuture` allows chaining of asynchronous computations without polling or the need of deeply nested callbacks (also known as callback hell). All `PgxFuture` instances returned by PGX APIs are instances of `CompletableFuture` and can be chained without the need of Java 8.

```

import java.util.concurrent.CompletableFuture

...

final GraphConfig graphConfig = ...
instance.createSessionAsync("my-session")
 .thenCompose(new Fun<PgxFuture, CompletableFuture<PgxFuture>>() {
 @Override
 public CompletableFuture<PgxFuture> apply(PgxFuture session) {
 return session.readGraphWithPropertiesAsync(graphConfig);
 }
 })
 .thenAccept(new Action<PgxFuture>() {
 @Override
 public void accept(PgxFuture graph) {
 // do something with loaded graph
 }
 });

```

The asynchronous chaining in the preceding example is explained as follows:

- The first line in the code makes an asynchronous call to `createSessionAsync()` to create a session. Once the promise is resolved, it returns a `PgxFuture` object, which is the newly created `PgxSession`.
- The code then calls the `.thenCompose()` handler by passing a function which takes the `PgxSession` object as an argument. Inside the function, there is another asynchronous `readGraphWithPropertiesAsync()` request which return another `PgxFuture` object. The outer `PgxFuture` object returned by `.thenCompose()` gets resolved when the `readGraphWithPropertiesAsync()` request completes.
- This is followed by the `.thenAccept()` handler. The function that is passed to `.thenAccept()` does not return anything. Therefore, the `future` return type of `.thenAccept()` is `PgxFuture<Void>`.

### Blocking Versus Chaining

For most use cases, you can block the caller thread. However, blocking can quickly lead to poor performance or deadlocks once things get more complex. As a rule, use blocking to quickly analyze selected graphs in a sequential manner, for example, in shell scripts or during interactive analysis using the interactive PGX shell.

Use chaining for applications built on top of PGX.

## 14.3.3 Cancelling Operation

You can cancel a pending request by invoking the `cancel` method of the returned `PgxFuture` instance.

For example:

```
PgxFuture<Object> promise=...
// do something else
promise.cancel(); // will cancel computation
```

Any subsequent calls to `promise.get()` will result in a `CancellationException` being thrown.

#### Note:

Due to Java's cooperative threading model, it might take some time before PGX actually stops the computation.

## 14.3.4 Handling Concurrent Asynchronous Operations

Using the `PgxSession#runConcurrently` API provided by the in-memory graph server (PGX), you can submit a list of suppliers of asynchronous APIs to run concurrently in the PGX server.

For example:

```
import oracle.pgx.api.*;

 Supplier<PgxFuture<?>> asyncRequest1 = () ->
session.readGraphWithPropertiesAsync(...);
 Supplier<PgxFuture<?>> asyncRequest2 = () ->
session.getAvailableSnapshotsAsync(...);

 List<Supplier<PgxFuture<?>>> supplierList = Arrays.asList(asyncRequest1,
asyncRequest2);

 //executing the async requests with the enabled optimization feature
List<?> results = session.runConcurrently(supplierList);

 //the supplied requests are mapped to their results and orderly collected
PgxGraph graph = (PgxGraph) results.get(0);
Deque<GraphMetaData> metaData = (Deque<GraphMetaData>) results.get(1);
```

## 14.4 Graph Client Sessions

The graph server (PGX) assumes there may be multiple concurrent clients, and each client submits request to the shared PGX server independently.

Each session has its own workspace in PGX and is isolated from other sessions.

You can share graphs or properties among sessions.

### Creating Sessions

The following methods in the `ServerInstance` class are used to create sessions:

- 
- [Java](#)
  - [Python](#)

### Java

```
PgxSession createSession(String source)
PgxSession createSession(String source, long idleTimeout, long taskTimeout,
TimeUnit unit)
```

The preceding methods accept the following arguments:

- `source` is any arbitrary string that describes the client. Currently, this string is only used for logging purposes.
- The user can specify the idle timeout (`idleTimeout`) and task timeout (`taskTimeout`) when creating a new session. If these values are not specified, default values are used. See [Configuration Parameters for the Graph Server \(PGX\) Engine](#) for more information on graph server (PGX) configuration options.

---

## Python

```
import pypgx
session = pypgx.get_session()
```

---

### Destroying Sessions

To destroy a session, simply call:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
session.destroyAsync();
```

## Java

```
session.destroy();
```

## Python

```
session.destroy()
```

---

Administrators can destroy sessions by ID using the following code:

```
instance.killSession(sessionId);
```

**Note:**

Calling administrative methods by default requires special authorization in client/server mode.

When a session is destroyed, PGX reclaims all of the resources associated with the session. Specifically, all transient data is destroyed immediately. See [Managing Transient Data](#) for more information on transient data.

However, PGX may choose to keep the loaded graph instance in memory for caching purposes, especially if a graph instance is shared by multiple clients. In summary, every graph remains in memory until no client is using it.

**Note:**

A session can be destroyed automatically via the session time-out mechanism. See [Configuration Parameters for the Graph Server \(PGX\) Engine](#) for more information on graph server (PGX) configuration options.

## 14.5 Graph Mutation and Subgraphs

This guide discusses the several methods provided by the graph server (PGX) for mutating graph instances.

You can use the mutation and subgraph methods that are defined in the `PgxGraph` class, to mutate a graph.

**Note:**

All of the mutating methods create a new graph or snapshot instance as the mutated version of the original graph, rather than mutating the original graph directly.

- [Altering Graphs](#)
- [Simplifying and Copying Graphs](#)
- [Transposing Graphs](#)
- [Undirecting Graphs](#)
- [Advanced Multi-Edge Handling](#)
- [Creating a Subgraph](#)
- [Creating a Bipartite Subgraph](#)
- [Creating a Sparsified Subgraph](#)

### 14.5.1 Altering Graphs

This section explains the graph alteration mutation used to add or remove vertex and edge providers of a graph.

You can add or remove vertex and edge providers in a graph that has been loaded or created previously. Providers can be added from existing datasources, or new empty providers can be created. The mutation can either create a new independent graph, or create a new snapshot for the graph.

The following topics explain in detail on adding and removing vertex and edge providers:

You must first create a graph-alteration builder to start altering an existing graph. For example, the following code shows how to start a graph alteration on a graph that is stored in a variable `graph`:



- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg-jshell> var alterationBuilder = graph.alterGraph()
```

## Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.graphalteration.GraphAlterationBuilder;

GraphAlterationBuilder alterationBuilder = graph.alterGraph();
```

## Python

```
alteration_builder = graph.alter_graph()
```

- 
- [Loading Or Removing Additional Vertex or Edge Providers](#)

### 14.5.1.1 Loading Or Removing Additional Vertex or Edge Providers

You can alter your graph by adding or removing vertex or edge providers from a specific datasource. Alternatively you can also add empty vertex or edge providers.

#### Keys in Additionally Loaded Providers

The vertex and edge providers that are loaded must provide the respective keys in accordance with the vertex ID and edge ID strategy of the graph being altered. If the ID strategy is `KEYS_AS_IDS`, the provider must create a key mapping. But, if the ID strategy is `UNSTABLE_GENERATED_IDS`, it must not create the key mapping.

- [Loading Vertex Providers](#)
- [Loading Edge Providers](#)
- [Adding Additional Empty Vertex or Edge Providers](#)
- [Removing Vertex or Edge Providers](#)
- [Applying the Alteration and Building a Graph or Snapshot](#)

#### 14.5.1.1.1 Loading Vertex Providers

You can add a vertex provider by calling `alterationBuilder.addVertexProvider(EntityProviderConfig vertexProviderConfig)`.

`vertexProviderConfig` is a vertex provider configuration and it provides configuration details such as:

- location of the datasource to load from
- the stored format
- properties of the vertex provider

### Adding a Vertex Provider from a JSON Configuration

You can add the provider by calling `alterationBuilder.addVertexProvider(String pathToVertexProviderConfig)` where `pathToVertexProviderConfig` points to a file accessible from the client that contains a JSON representation of a vertex provider configuration.

For example, a vertex provider configuration can be stored in a JSON file as shown:

```
{
 "name": "Accounts",
 "format": "rdbms",
 "database_table_name": "BANK_ACCOUNTS",
 "key_column": "ID",
 "key_type": "integer",
 "props": [
 {
 "name": "ID",
 "type": "integer"
 },
 {
 "name": "NAME",
 "type": "string"
 }
]
}
```

You can then add the vertex provider as shown in the following example:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
// Loading by indicating the path to the JSON file
opg4j> alterationBuilder.addVertexProvider("<path-to-vertex-provider-configuration>")
$9 ==>
oracle.pgx.api.graphalteration.internal.GraphAlterationBuilderImpl@48d464cf

// Or by first loading the content of a JSON file into an
EntityProviderConfig object
opg4j> EntityProviderConfig vertexProviderConfig = new
AnyFormatEntityProviderConfigFactory().fromPath("<path-to-vertex-provider-configuration>")
```

```
vertexProviderConfig ==>
{"format":"rdbms","name":"Accounts","database_table_name":"BANK_ACCOUNTS",
"loading":{"create_key_mapping":true},"key_type":"integer","props":
[{"type":"integer","name":"ID"},
{"type":"string","name":"NAME"}],"key_column":"ID"}
opg4j> alterationBuilder.addVertexProvider(vertexProviderConfig)
$15 ==>
oracle.pgx.api.graphalteration.internal.GraphAlterationBuilderImpl@77e2a5d3
```

## Java

```
// Loading by indicating the path to the JSON file
alterationBuilder.addVertexProvider("<path-to-vertex-provider-configuration>");

// Or by first loading the content of a JSON file into an
EntityProviderConfig object
EntityProviderConfig vertexProviderConfig = new
AnyFormatEntityProviderConfigFactory().fromPath("<path-to-vertex-provider-configuration>");
alterationBuilder.addVertexProvider(vertexProviderConfig);
```

## Python

```
Loading by indicating the path to the JSON file
alterationBuilder.add_vertex_provider("<path-to-vertex-provider-configuration>");
```

---

### Adding a Vertex Provider Programmatically Using an API

Alternatively, the vertex provider configuration can be built programmatically:

- [JShell](#)
- [Java](#)

## JShell

```
opg4j> RdbmsEntityProviderConfigBuilder vertexProviderConfigBuilder =
new RdbmsEntityProviderConfigBuilder().
...> setName("Accounts").
...> setKeyColumn("ID").
...>
setDatabaseTableName("BANK_ACCOUNTS").
...> addProperty("ID",
PropertyType.INTEGER)
vertexProviderConfigBuilder ==>
```

```
oracle.pgx.config.RdbmsEntityProviderConfigBuilder@8ff4d2b

opg4j> EntityProviderConfig vertexProviderConfig =
vertexProviderConfigBuilder.build()
vertexProviderConfig ==> {"error_handling":
{"format":"rdbms","name":"Accounts","database_table_name":"BANK_ACCOUNTS",
loading":{"create_key_mapping":true},"attributes":
{"key_type":"long","props":
[{"dimension":0,"type":"integer","name":"ID"}],"key_column":"ID"}

opg4j> alterationBuilder.addVertexProvider(vertexProviderConfig)
$24 ==>
oracle.pgx.api.graphalteration.internal.GraphAlterationBuilderImpl@7b303608
```

## Java

```
RdbmsEntityProviderConfigBuilder vertexProviderConfigBuilder = new
RdbmsEntityProviderConfigBuilder()
 .setName("Accounts")
 .setKeyColumn("ID")
 .setDatabaseTableName("BANK_ACCOUNTS")
 .addProperty("ID", PropertyType.INTEGER);
EntityProviderConfig vertexProviderConfig =
vertexProviderConfigBuilder.build();
alterationBuilder.addVertexProvider(vertexProviderConfig);
```

### 14.5.1.1.2 Loading Edge Providers

You can add an edge provider by calling `alterationBuilder.addEdgeProvider(EntityProviderConfig edgeProviderConfig)` where `edgeProviderConfig`. `edgeProviderConfig` is an edge provider configuration and it provides configuration details such as:

- location of the datasource to load from
- the stored format
- properties of the edge provider

The source and destination vertex providers to which it is linked must either be already in the base graph (and not removed in the alteration), or added with the alteration.

#### Adding an Edge Provider from a JSON Configuration

You can also add the provider by calling `alterationBuilder.addEdgeProvider(String pathToEdgeProviderConfig)` where `pathToEdgeProviderConfig` points to a file accessible from the client that contains a JSON representation of an edge provider configuration.

For example, an edge provider configuration can be stored in a JSON file as shown:

```
{
 "name": "Transfers",
 "format": "rdbms",
```

```

 "database_table_name": "BANK_EDGES_AMT",
 "key_column": "ID",
 "source_column": "SRC_ID",
 "destination_column": "DEST_ID",
 "source_vertex_provider": "Accounts",
 "destination_vertex_provider": "Accounts",
 "props": [
 {
 "name": "AMOUNT",
 "type": "float"
 }
]
 }
}

```

You can then add the edge provider as shown in the following example:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```

// Loading by indicating the path to the JSON file
opg4j> alterationBuilder.addEdgeProvider("<path-to-edge-provider-
configuration>")
$10 ==>
oracle.pgx.api.graphalteration.internal.GraphAlterationBuilderImpl@48d4
64cf

// Or by first loading the content of a JSON file into an
EntityProviderConfig object
opg4j> EntityProviderConfig edgeProviderConfig = new
AnyFormatEntityProviderConfigFactory().fromPath("<path-to-edge-
provider-configuration>")
edgeProviderConfig ==>
{"format":"rdbms","source_vertex_provider":"Accounts","name":"Transfers
","database_table_name":"BANK_EDGES_AMT","loading":
{"create_key_mapping":false},"source_column":"SRC_ID","destination_colu
mn":
"DEST_ID","key_type":"long","destination_vertex_provider":"Accounts","p
rops":[{"type":"float","name":"AMOUNT"}],"key_column":"ID"}

opg4j> alterationBuilder.addEdgeProvider(edgeProviderConfig)
$26 ==>
oracle.pgx.api.graphalteration.internal.GraphAlterationBuilderImpl@7b30
3608

```

## Java

```
// Loading by indicating the path to the JSON file
alterationBuilder.addEdgeProvider("<path-to-edge-provider-configuration>");

// Or by first loading the content of a JSON file into an
EntityProviderConfig object
EntityProviderConfig edgeProviderConfig = new
AnyFormatEntityProviderConfigFactory().fromPath("<path-to-edge-provider-
configuration>");
alterationBuilder.addEdgeProvider(edgeProviderConfig);
```

## Python

```
Loading by indicating the path to the JSON file
alterationBuilder.add_edge_provider("<path-to-edge-provider-configuration>");
```

---

### Adding an Edge Provider Programmatically Using an API

Alternatively, the edge provider configuration can be built programmatically:

- [JShell](#)
- [Java](#)

### JShell

```
opg4j> RdbmsEntityProviderConfigBuilder edgeProviderConfigBuilder = new
RdbmsEntityProviderConfigBuilder().
...> setName("Transfers").
...> setKeyColumn("id").
...> setSourceColumn("src_id").
...>
setDestinationColumn("dest_id").
...>
setSourceVertexProvider("Accounts").
...>
setDestinationVertexProvider("Accounts").
...> createKeyMapping(true).
...>
setDatabaseTableName("bank_txns").
...>
addProperty("from_acct_id", PropertyType.LONG).
...> addProperty("to_acct_id",
PropertyType.LONG).
...> addProperty("amount",
PropertyType.LONG)
edgeProviderConfigBuilder ==>
oracle.pgx.config.RdbmsEntityProviderConfigBuilder@5a5f65b9
```

```

opg4j> EntityProviderConfig edgeProviderConfig =
edgeProviderConfigBuilder.build()
edgeProviderConfig ==> {"error_handling":
{"attributes{}", "destination_column": "dest_id", "key_type": "long", "dest
ination_vertex_provider": "Accounts", "key_column": "id", "format": "rdbms",
"source_vertex_provider":
"Accounts", "name": "Transfers", "database_table_name": "bank_txns", "loadin
g": {"create_key_mapping": true}, "source_column": "src_id", "props":
[{"dimension": 0, "type": "long", "name": "from_acct_id"},
{"dimension": 0, "type": "long",
"name": "to_acct_id"}, {"dimension": 0, "type": "long", "name": "amount"}]}

opg4j> alterationBuilder.addEdgeProvider(edgeProviderConfig)
$30 ==>
oracle.pgx.api.graphalteration.internal.GraphAlterationBuilderImpl@441c
cfd7

```

## Java

```

RdbmsEntityProviderConfigBuilder edgeProviderConfigBuilder = new
RdbmsEntityProviderConfigBuilder()
.setName("Transfers")
.setKeyColumn("id")
.setSourceColumn("src_id")
.setDestinationColumn("dest_id").
.setSourceVertexProvider("Accounts")
.setDestinationVertexProvider("Accounts")
.createKeyMapping(true)
.setDatabaseTableName("bank_txns")
.addProperty("from_acct_id", PropertyType.LONG)
.addProperty("to_acct_id", PropertyType.LONG)
.addProperty("amount", PropertyType.LONG);

EntityProviderConfig edgeProviderConfig =
edgeProviderConfigBuilder.build();
alterationBuilder.addEdgeProvider(edgeProviderConfig);

```

### 14.5.1.1.3 Adding Additional Empty Vertex or Edge Providers

You can also add empty vertex or edge providers, without having the providers connected to any specific datasource.

The names and types of the properties of each empty provider can be specified programmatically. Similarly, you can also specify if a key mapping for the providers needs to be created.

#### Adding Additional Empty Vertex Providers

You can add an empty vertex provider by calling `alterationBuilder.addEmptyVertexProvider(String vertexProviderName)`. You

can then add properties, specify the key column, create the key mapping programmatically as shown in the following example.

See the **GraphAlterationEmptyVertexProviderBuilder** Interface in the Javadoc for more details.

- [JShell](#)
- [Java](#)

## JShell

```
opg4j> alterationBuilder.addEmptyVertexProvider("AccountsProvider").
...> setLabel("Accounts").
...> createKeyMapping(true).
...> addProperty("NAME", PropertyType.STRING)
$14 ==>
oracle.pgx.api.graphalteration.internal.GraphAlterationEmptyVertexProviderBui
lderImpl@4b3ea082
```

## Java

```
alterationBuilder.addEmptyVertexProvider("AccountsProvider")
 .setLabel("Accounts")
 .createKeyMapping(true)
 .addProperty("NAME", PropertyType.STRING);
```

---

### Adding Additional Empty Edge Providers

You can add an empty edge provider by calling `alterationBuilder.addEmptyEdgeProvider(String providerName, String sourceProvider, String destProvider)`. You can then add properties, specify the key column, create the key mapping programmatically as shown in the following example.

See the **GraphAlterationEmptyEdgeProviderBuilder** Interface in the Javadoc for more details.

- [JShell](#)
- [Java](#)

## JShell

```
opg4j> alterationBuilder.addEmptyEdgeProvider("TransactionProvider",
"Accounts", "Accounts").
...> setLabel("Transfers").
...> createKeyMapping(false). // set to false if no keys are needed
```



```
...> addProperty("Description", PropertyType.STRING)
$26 ==>
oracle.pgx.api.graphalteration.internal.GraphAlterationEmptyEdgeProvide
rBuilderImpl@54720caf
```

## Java

```
alterationBuilder.addEmptyEdgeProvider("TransactionProvider",
"Accounts", "Accounts")
 .setLabel("Transfers")
 .createKeyMapping(false)
 .addProperty("Description", PropertyType.STRING);
```

---

### 14.5.1.1.4 Removing Vertex or Edge Providers

You can remove an edge provider by calling `alterationBuilder.removeEdgeProvider(String edgeProviderName)`, where `edgeProviderName` is the name of the edge provider to be removed from the graph.

Similarly, calling `alterationBuilder.removeVertexProvider(String vertexProviderName)` will result in the graph to not contain that specific vertex provider. If that vertex provider was the source or destination provider for some edge providers in the base graph, those edge providers should also be removed before the application of the alteration or an exception will be thrown.

It is possible to indicate that the edge providers associated to a removed vertex provider should be automatically removed by calling `alterationBuilder.cascadeEdgeProviderRemovals(boolean cascadeEdgeProviderRemovals)` with `cascadeEdgeProviderRemovals` set to `true`.

### 14.5.1.1.5 Applying the Alteration and Building a Graph or Snapshot

You must call `alterationBuilder.build()`, once the different vertex and edge providers have been added or removed in the alteration to actually apply the operation. By calling `alterationBuilder.build()`, a new graph is created and that graph contains all the providers of the base graph excluding the removed providers, and the additionally loaded providers.

You can also call `alterationBuilder.buildNewSnapshot()`, in which case, a new snapshot for the base graph is created and that snapshot contains all the providers of the base graph excluding the removed providers, and the additionally loaded providers.

## 14.5.2 Simplifying and Copying Graphs

You can create a simplified version of the graph by calling the `simplify()` method.

- 
- [Java](#)

- [Python](#)

## Java

```
PgxGraph simplify(Collection<VertexProperty<?, ?>> vertexProps,
 Collection<EdgeProperty<?>> edgeProps, MultiEdges multiEdges,
 SelfEdges selfEdges, TrivialVertices trivialVertices,
 Mode mode, String newGraphName)
```

## Python

```
simplify(self, vertex_properties=True, edge_properties=True, keep_multi_edges=False,
 keep_self_edges=False, keep_trivial_vertices=False, in_place=False,
 name=None)
```

---

The first two arguments (`vertexProps` and `edgeProps`) list which properties will be copied into the newly created simplified graph instance. PGX provides convenience constants `VertexProperty.ALL`, `EdgeProperty.ALL` and `VertexProperty.NONE`, `EdgeProperty.NONE` to specify all properties or none properties to be stored, respectively.

The next three arguments determine which operations will be performed to simplify the graph.

- `multiEdges`: if `MultiEdges.REMOVE_MULTI_EDGES`, eliminate multiple edges between a source vertex and a destination vertex, that is, leave at most one edge between two vertices. `MultiEdges.KEEP_MULTI_EDGES` indicates to keep them. By default, PGX picks one edge out of the multi-edges and takes its properties. See [Advanced Multi-Edge Handling](#) for more fine-grained control over the edge properties during simplification.
- `selfEdges`: if `SelfEdges.REMOVE_SELF_EDGES`, eliminate every edge whose source and destination are the same vertex. `SelfEdges.KEEP_MULTI_EDGES` indicates to keep them.
- `trivialVertices`: if `TrivialVertices.REMOVE_TRIVIAL_VERTICES`, eliminate all the vertices that have neither incoming edges nor outgoing edges. `TrivialVertices.KEEP_TRIVIAL_VERTICES` indicates to keep them.

The `mode` argument, if set to `Mode.MUTATE_IN_PLACE`, requests that the mutation occurs directly on the specified graph instance without creating a new one. If set to `Mode.CREATE_COPY`, the method will create a new graph instance with the new name in `newGraphName`. If `newGraphName` is omitted (or null), PGX will generate a unique graph name.

The return value of this method is the simplified `PgxGraph` instance.

The `Mode.MUTATE_IN_PLACE` option is only applicable if the graph is marked as mutable. Every graph is immutable by default when loaded into PGX. To make a `PgxGraph` mutable, the client should create a private copy of the graph first, using one of the following methods:

- 
- [Java](#)
  - [Python](#)

## Java

```
PgxGraph clone()
PgxGraph clone(String newGraphName)
PgxGraph clone(Collection<VertexProperty<?, ?>> vertexProps,
Collection<EdgeProperty<?>> edgeProps, String newGraphName)
```

## Python

```
clone(self, vertex_properties=True, edge_properties=True, name=None)
```

---

As with `simplify()`, the user can specify optional properties of the graph to copy with `vertexProps` and `edgeProps`. If no properties are specified, all of the original graph's properties will be copied into the new graph instance. The user can specify the name of the newly created graph instance with `newGraphName`.

## 14.5.3 Transposing Graphs

You can create a transposed version of the graph.

- 
- [Java](#)
  - [Python](#)

## Java

```
PgxGraph transpose(Collection<VertexProperty<?, ?>> vertexProps,
Collection<EdgeProperty<?>> edgeProps,
Map<String, String> edgeLabelMapping,
Mode mode, String newGraphName)
```

## Python

```
transpose(self, vertex_properties=True, edge_properties=True,
edge_label_mapping=None, in_place=False,
name=None)
```

---

The `edgeLabelMapping` argument can be used to rename edge labels. If any key in the given map does not exist as an edge label, it will be ignored.

`edgeLabelMapping` argument can also be an empty Map or null.

- `null`: if argument is `null`, edge labels from source graph will be removed on transposed graph. (default behavior when using convenience methods).

- `empty Map`: if argument is an empty Map, edge labels from source graph will be neither removed or renamed. Instead, it will be kept as it is in source graph.

See [Simplifying and Copying Graphs](#) for the meaning of the other parameters.

Additionally, the graph server (PGX) provides the following convenience methods from the `PgxGraph` class for the common operation of copying all vertex and edge properties into the transposed graph instance:

- `transpose(Mode mode, String newGraphName)`
- `transpose(String newGraphName)`
- `transpose(Mode mode)`

## 14.5.4 Undirecting Graphs

The following methods create the undirected version of a graph instance:

- 
- [Java](#)
  - [Python](#)

### Java

```
PgxGraph undirect()
PgxGraph undirect(String newGraphName)
PgxGraph undirect(MultiEdges multiEdges, SelfEdges selfEdges,
TrivialVertices trivialVertices, Mode mode, String newGraphName)
PgxGraph undirect(Collection<VertexProperty<?, ?>> vertexProps,
Collection<EdgeProperty<?>> edgeProps, MultiEdges multiEdges, SelfEdges
selfEdges, Mode mode, String newGraphName)
```

### Python

```
undirect(self, vertex_properties=True, edge_properties=True,
keep_multi_edges=True, keep_self_edges=True,
keep_trivial_vertices=True, in_place=False, name=None)
```

---

The first two methods create an undirected version of the graph while copying all of the vertex properties. `newGraphName` is an optional argument to specify the name of the newly created graph instance.

In contrast, the third and fourth methods concurrently perform *undirecting* and *simplifying* of a graph. See [Simplifying and Copying Graphs](#) for the meaning of each parameter.

All methods return an object of the undirected `PgxGraph` type.

An undirected graph has some restrictions. Some algorithms are only supported on directed graphs or are not yet supported for undirected graphs. Further, PGX does not support to store undirected graphs nor reading from undirected formats. Since the edges do not have a

direction anymore, the behavior of `pgxEdge.getSource()` or `pgxEdge.getDestination()` can be ambiguous. In order to provide deterministic results, PGX will always return the vertex with the smaller internal id as source and the other as destination vertex.

## 14.5.5 Advanced Multi-Edge Handling

Both `simplify()` and `undirect()` support the removal of multi-edges using `MultiEdges.REMOVE_MULTI_EDGES`. If this parameter is set, all multi-edges in this graph are removed, that is, collapsed. Whenever several multi-edges with edge properties are collapsed into one edge, you can choose one of the following two strategies supported by the graph server (PGX) to decide how to treat the corresponding properties:

- Picking
- Merging

If you choose picking, the graph server (PGX) picks one edge out of every set of multi-edges and copies all its properties including the edge label and key into the new graph. In the case of merging, the graph server (PGX) creates a completely new edge out for every set of multi-edges. PGX determines the properties of these new edges by applying a `MergingFunction` on every property of the multi-edges.

If there are no multi-edges between two vertices, that is, zero or only one edge, the chosen strategy does not have an effect on the outcome. The edge is kept with all its properties as it is.

- [Picking](#)
- [Merging](#)
- [StrategyBuilder in General](#)

### 14.5.5.1 Picking

This strategy can be used to pick an edge out of multi-edges. The graph server (PGX) allows the user to define several picking criteria. You can pick by:

- Property
- Label
- Edge-ID

Every picking criteria has to be combined with a `PickingStrategyFunction`. PGX supports either `PickingStrategyFunction.MIN` and `PickingStrategyFunction.MAX`, which picks the edge whose property/label/id is either minimal or maximal. If one does not specify a picking criteria, PGX will non-deterministically pick an edge out of the multi-edges.

A `PickingStrategy` can be created using a `PickingStrategyBuilder`, which can be retrieved by calling `createPickingStrategyBuilder()` on the target graph.


You can call one of the following functions as per your chosen picking criteria:

```
PickingStrategyBuilder setPickByEdgeId(PickingStrategyFunction
pickingStrategyFunction)
PickingStrategyBuilder setPickByLabel(PickingStrategyFunction
pickingStrategyFunction)
```

```
PickingStrategyBuilder setPickByProperty(EdgeProperty edgeProperty,
PickingStrategyFunction pickingStrategyFunction)
PickingStrategyBuilder setPickByProperty(String propertyName, PickingStrategyFunction
pickingStrategyFunction)
```

The following figure shows how PGX picks the edge with the *minimal* cost and takes all its properties.

**Figure 14-1 Picking Strategy**



### 14.5.5.2 Merging

This strategy can be used to merge the properties of multi-edges. The graph server (PGX) allows the user to define a `MergingFunction` for every property. Currently, PGX supports the following functions:


- `MergingFunction.MIN`
- `MergingFunction.MAX`
- `MergingFunction.SUM`

 **Note:**

SUM is only defined on numeric properties.

The following figure shows how the graph server (PGX) merges the different edge properties and labels. It takes the *minimal* cost, the *sum* of distances and the *maximal* edge label.

**Figure 14-2 Merging Strategy**



### 14.5.5.3 StrategyBuilder in General

By default, both the StrategyBuilders use the same values as in the convenience methods of `simplify()` and `undirect()`. This includes that all properties are kept by default. If one wants to drop specific properties, one can either use the `dropVertexProperty()` or `dropEdgeProperty()` functions.

```
MutationStrategyBuilder setNewGraphName(String newGraphName)
MutationStrategyBuilder setCopyMode(Mode mode)
MutationStrategyBuilder setTrivialVertices(TrivialVertices
trivialVertices)
MutationStrategyBuilder setSelfEdges(SelfEdges selfEdges)
MutationStrategyBuilder setMultiEdges(MultiEdges multiEdges)
MutationStrategyBuilder
dropVertexProperties(Collection<VertexProperty<?, ?>> vertexProperty)
MutationStrategyBuilder dropEdgeProperties(Collection<EdgeProperty<?>>
edgeProperty)
MutationStrategyBuilder dropVertexProperty(VertexProperty<?, ?>
vertexProperty)
MutationStrategyBuilder dropEdgeProperty(EdgeProperty<?> edgeProperty)
MutationStrategyBuilder build()
```

`Simplify()` and `undirect()` can be called using a `MutationStrategy` as follows:

```
MutationStrategy strategy = strategyBuilder.build()
PgxGraph simplifiedGraph graph.simplify(strategy)
//OR
PgxGraph undirectedGraph graph.undirect(strategy)
```

## 14.5.6 Creating a Subgraph

PGX provides the following methods for creating subgraphs via a filter (see [Filter Expressions](#) for more information) expression:

- 
- [Java](#)
  - [Python](#)

### Java

```
PgxGraph filter(GraphFilter graphFilter)
PgxGraph filter(GraphFilter graphFilter, String newGraphName)
PgxGraph filter(Collection<VertexProperty<?, ?>> vertexProps,
Collection<EdgeProperty<?>> edgeProps, GraphFilter graphFilter, String
newGraphName)
```

### Python

```
filter(self, graph_filter, vertex_properties=True, edge_properties=True,
name=None)
```

---

As in the other graph mutating methods, the user has the option to specify the name of the subgraph with the `newGraphName` parameter and of choosing the vertex and edge properties to be copied into the subgraph (`vertexProps` and `edgeProps`). All of the preceding methods return a `PgxGraph` object which represents the created subgraph.

All filter methods require a `GraphFilter` argument containing a filter expression. Fundamentally, the filter expression is a Boolean expression that is evaluated for every vertex and edge in the original graph (in parallel). If the expression is evaluated as `true` for the vertex or edge, then that vertex or edge is included in the subgraph.

See [Creating Subgraphs](#) for more information on how to create subgraphs from graphs loaded into memory.

## 14.5.7 Creating a Bipartite Subgraph

The graph server (PGX) enables the client to create a bipartite subgraph. The following methods return the created `BipartiteGraph` instance:



- [Java](#)
- [Python](#)

## Java

```
BipartiteGraph bipartiteSubGraphFromLeftSet(VertexSet<?> vertexSet)
BipartiteGraph bipartiteSubGraphFromLeftSet(VertexSet<?> vertexSet,
String newGraphName)
BipartiteGraph
bipartiteSubGraphFromLeftSet(Collection<VertexProperty<?, ?>>
vertexProps, Collection<EdgeProperty<?>> edgeProps, VertexSet<?>
vertexSet, String newGraphName)
BipartiteGraph
bipartiteSubGraphFromLeftSet(Collection<VertexProperty<?, ?>>
vertexProps, Collection<EdgeProperty<?>> edgeProps, VertexSet<?>
vertexSet, String newGraphName, String isLeftPropName)
```

## Python

```
bipartite_sub_graph_from_left_set(self, vset, vertex_properties=True,
edge_properties=True, name=None, is_left_name=None)
```

These methods require an additional argument `vertexSet`, which points to a set of vertices (see [Using Collections and Maps](#) for more information) whose elements (vertices) would contain the left vertices (that is, vertices on the left side of the bipartite graph that have only edges to vertices on the right side) in the resulting bipartite graph.

When creating the bipartite subgraph, PGX automatically inserts an additional `boolean` vertex property `isLeft`. The value of this property is set `true` for the left vertices and `false` for the right vertices in the bipartite subgraph. The name of the `isLeft` vertex property can be obtained with `getIsLeftPropertyAsync()` on the returned `BipartiteGraph` object.

The user has the option to specify a name for the newly created graph (`newGraphName`) as well as a custom name for the Boolean left-vertex indicating property (`isLeftPropName`). The user can also specify the vertex and edge properties to be copied into the newly created graph instance (`vertexProps` and `edgeProps`).

## 14.5.8 Creating a Sparsified Subgraph

The graph server (PGX) supports creating a sparsified subgraph of a graph:

- [Java](#)
- [Python](#)

## Java

```
PgxGraph sparsify(double e)
PgxGraph sparsify(double e, String newGraphName)
PgxGraph sparsify(Collection<VertexProperty<?, ?>> vertexProps,
Collection<EdgeProperty<?>> edgeProps, double e, String newGraphName)
```

## Python

```
sparsify(self, sparsification, vertex_properties=True, edge_properties=True,
name=None)
```

---

The `double` argument `e` is the sparsification coefficient with a value between 0.0 and 1.0.

The user again has the option to specify the name for the newly created graph (`newGraphName`) as well as the vertex and edge properties to be copied into the newly created graph instance (`vertexProps` and `edgeProps`).

The returned `PgxGraph` object represents a sparsified subgraph which has fewer edges than the original graph.

## 14.6 Graph Builder and Graph Change Set

This guide explains the `GraphBuilder` API used for creating graphs and the `GraphChangeSet` interface used for modifying loaded graphs.

- [Building Graphs Using GraphBuilder Interface](#)
- [Modifying Loaded Graphs Using ChangeSet](#)

### 14.6.1 Building Graphs Using GraphBuilder Interface

Using the `GraphBuilder` interface, you can create graphs programmatically.

The basic work flow for creating graphs from scratch is:

1. Acquire a modifiable graph builder to accumulate all the new vertices and edges
2. Add vertices and edges to the graph builder
3. Create a `PgxGraph` out of the accumulated changes

- [Creating a Simple Graph](#)
- [Adding a Vertex Property](#)
- [Using Strings as Vertex Identifiers](#)
- [Referencing a Vertex for Creating Edges](#)
- [Adding an Edge Property and a Label](#)
- [Using Graph Builder with Implicit IDs](#)

## 14.6.1.1 Creating a Simple Graph

This section shows an example of creating a simple graph using the `createGraphBuilder()` method.

- [JShell](#)
- [Java](#)
- [Python](#)

### JShell

```
opg4j> var builder = session.createGraphBuilder()
builder ==> GraphBuilderImpl[session=cd201ac9-e73f-447c-9cec-
cd929293acc3,vertexChanges=0,edgeChanges=0]

opg4j> builder.addEdge(1, 2)
opg4j> builder.addEdge(2, 3)
opg4j> builder.addEdge(2, 4)
opg4j> builder.addEdge(3, 4)
opg4j> builder.addEdge(4, 2)

opg4j> var graph = builder.build()
graph ==>
PgxGraph[name=anonymous_graph_16,N=4,E=5,created=1629805890550]
```

### Java

```
import oracle.pgx.api.*;

PgxSession session = Pgx.createSession("example");
GraphBuilder<Integer> builder = session.createGraphBuilder();

builder.addEdge(1, 2);
builder.addEdge(2, 3);
builder.addEdge(2, 4);
builder.addEdge(3, 4);
builder.addEdge(4, 2);

PgxGraph graph = builder.build();
```

### Python

```
from pypgx import get_session

session = get_session(session_name="example")
builder = session.create_graph_builder()
builder.add_edge(1, 2)
builder.add_edge(2, 3)
```

```
builder.add_edge(2, 4)
builder.add_edge(3, 4)
builder.add_edge(4, 2)
graph = builder.build()
```

---

Also, note that the following:

- A call to `addEdge` consists of the new unique edge ID, the source vertex ID and the destination vertex ID.
- No graph configuration is required.
- When adding edges, all vertices that do not already exist are created on the fly as edges are created.
- `GraphBuilder` supports only the following two generation strategies for creating vertices and edge IDs:
  - `USER_IDS` (the default value)
  - `AUTO_GENERATED`

### 14.6.1.2 Adding a Vertex Property

You can also add vertices separately and assign property values to them.

The following example shows how to add a vertex property using the `GraphBuilder` interface.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> var builder = session.createGraphBuilder()
opg4j> builder.addVertex(1).setProperty("double-prop", 0.1)
opg4j> builder.addVertex(2).setProperty("double-prop", 2.0)
opg4j> builder.addVertex(3).setProperty("double-prop", 0.3)
opg4j> builder.addVertex(4).setProperty("double-prop", 4.56789)
opg4j> builder.addEdge(1, 2)
opg4j> builder.addEdge(2, 3)
opg4j> builder.addEdge(2, 4)
opg4j> builder.addEdge(3, 4)
opg4j> builder.addEdge(4, 2)
opg4j> var graph = builder.build()
```

#### Java

```
import oracle.pgx.api.*;
```

```

PgxSession session = Pgx.createSession("example");
GraphBuilder<Integer> builder = session.createGraphBuilder();

builder.addVertex(1).setProperty("double-prop", 0.1);
builder.addVertex(2).setProperty("double-prop", 2.0);
builder.addVertex(3).setProperty("double-prop", 0.3);
builder.addVertex(4).setProperty("double-prop", 4.56789);

builder.addEdge(1, 2);
builder.addEdge(2, 3);
builder.addEdge(2, 4);
builder.addEdge(3, 4);
builder.addEdge(4, 2);

PgxGraph graph = builder.build();

```

## Python

```

from pypgx import get_session

session = get_session(session_name="example")
builder = session.create_graph_builder()

builder.add_vertex(1).set_property("double-prop", 0.1)
builder.add_vertex(2).set_property("double-prop", 2.0)
builder.add_vertex(3).set_property("double-prop", 0.3)
builder.add_vertex(4).set_property("double-prop", 4.56789)

builder.add_edge(1, 2)
builder.add_edge(2, 3)
builder.add_edge(2, 4)
builder.add_edge(3, 4)
builder.add_edge(4, 2)

graph=builder.build()

```

If the value for a property is missing for a vertex or an edge, a default value is assumed as shown:

**Table 14-4 Default Property Values**

Properties	Default Values
Numeric	0 (or the respective equivalent)
Boolean	false
Date	1.1.1970 00:00:00
String	null

**Tip:**

Multiple calls to `setProperty` can be chained to set multiple property values at once.

### 14.6.1.3 Using Strings as Vertex Identifiers

By default, integer vertex IDs are used to identify a vertex. But, the type of the vertex ID can also be a long or a string.

In order to implement this, you must specify the vertex ID type when creating the graph using the `GraphBuilder` as shown:

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg4j> GraphBuilder<String> builder =
session.createGraphBuilder(IdType.STRING)
opg4j> builder.addVertex("vertex 1").setProperty("double-prop", 0.1)
opg4j> builder.addVertex("vertex 2").setProperty("double-prop", 2.0)
opg4j> builder.addVertex("vertex 3").setProperty("double-prop", 0.3)
opg4j> builder.addVertex("vertex 4").setProperty("double-prop", 4.56789)
opg4j> builder.addEdge("vertex 1", "vertex 2")
opg4j> builder.addEdge("vertex 2", "vertex 3")
opg4j> builder.addEdge("vertex 2", "vertex 4")
opg4j> builder.addEdge("vertex 3", "vertex 4")
opg4j> builder.addEdge("vertex 4", "vertex 2")

opg4j> var graph = builder.build()
```

#### Java

```
import oracle.pgx.api.*;
import oracle.pgx.common.types.IdType;

PgxSession session = Pgx.createSession("example");
GraphBuilder<String> builder = session.createGraphBuilder(IdType.STRING);

builder.addVertex("vertex 1").setProperty("double-prop", 0.1);
builder.addVertex("vertex 2").setProperty("double-prop", 2.0);
builder.addVertex("vertex 3").setProperty("double-prop", 0.3);
builder.addVertex("vertex 4").setProperty("double-prop", 4.56789);

builder.addEdge("vertex 1", "vertex 2");
builder.addEdge("vertex 2", "vertex 3");
```

```
builder.addEdge("vertex 2", "vertex 4");
builder.addEdge("vertex 3", "vertex 4");
builder.addEdge("vertex 4", "vertex 2");

PgxGraph graph = builder.build();
```

## Python

```
from pypgx import get_session

session = get_session(session_name="example")
builder = session.create_graph_builder(id_type='string')
builder.add_vertex("vertex 1").set_property("double-prop", 0.1)
builder.add_vertex("vertex 2").set_property("double-prop", 2.0)
builder.add_vertex("vertex 3").set_property("double-prop", 0.3)
builder.add_vertex("vertex 4").set_property("double-prop", 4.56789)

builder.add_edge("vertex 1", "vertex 2")
builder.add_edge("vertex 2", "vertex 3")
builder.add_edge("vertex 2", "vertex 4")
builder.add_edge("vertex 3", "vertex 4")
builder.add_edge("vertex 4", "vertex 2")

graph = builder.build()
```

---

### 14.6.1.4 Referencing a Vertex for Creating Edges

You can also avoid entering the full vertex ID when adding an edge. For this, you must obtain a reference to the vertex that is created, which can be later used in the `addEdge` statement.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> GraphBuilder<String> builder =
session.createGraphBuilder(IdType.STRING)

opg4j> var v1 = builder.addVertex("vertex 1").setProperty("double-
prop", 0.1)
opg4j> var v2 = builder.addVertex("vertex 2").setProperty("double-
prop", 2.0)
opg4j> var v3 = builder.addVertex("vertex 3").setProperty("double-
prop", 0.3)
opg4j> var v4 = builder.addVertex("vertex 4").setProperty("double-
```

```
prop", 4.56789)

opg4j> builder.addEdge(v1, v2)
opg4j> builder.addEdge(v2, v3)
opg4j> builder.addEdge(v2, v4)
opg4j> builder.addEdge(v3, v4)
opg4j> builder.addEdge(v4, v2)

opg4j> var graph = builder.build()
```

## Java

```
import oracle.pgx.api.*;
import oracle.pgx.common.types.IdType;

PgxSession session = Pgx.createSession("example");
GraphBuilder<String> builder = session.createGraphBuilder(IdType.STRING);

VertexBuilder<String> v1 = builder.addVertex("vertex 1").setProperty("double-
prop", 0.1);
VertexBuilder<String> v2 = builder.addVertex("vertex 2").setProperty("double-
prop", 2.0);
VertexBuilder<String> v3 = builder.addVertex("vertex 3").setProperty("double-
prop", 0.3);
VertexBuilder<String> v4 = builder.addVertex("vertex 4").setProperty("double-
prop", 4.56789);

builder.addEdge(v1, v2);
builder.addEdge(v2, v3);
builder.addEdge(v2, v4);
builder.addEdge(v3, v4);
builder.addEdge(v4, v2);

PgxGraph graph = builder.build();
```

## Python

```
from pypgx import get_session

session = get_session(session_name="example")
builder = session.create_graph_builder(id_type='string')

v1 = builder.add_vertex("vertex 1").set_property("double-prop", 0.1)
v2 = builder.add_vertex("vertex 2").set_property("double-prop", 2.0)
v3 = builder.add_vertex("vertex 3").set_property("double-prop", 0.3)
v4 = builder.add_vertex("vertex 4").set_property("double-prop", 4.56789)

builder.add_edge(v1, v2)
builder.add_edge(v2, v3)
builder.add_edge(v2, v4)
builder.add_edge(v3, v4)
builder.add_edge(v4, v2)
```



```
graph = builder.build()
```

---

### 14.6.1.5 Adding an Edge Property and a Label

The following examples show how to add an edge property and a label to a graph.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> var builder = session.createGraphBuilder(IdType.STRING)

opg4j> var v1 = builder.addVertex("vertex 1").setProperty("double-
prop", 0.1)
opg4j> var v2 = builder.addVertex("vertex 2").setProperty("double-
prop", 2.0)
opg4j> var v3 = builder.addVertex("vertex 3").setProperty("double-
prop", 0.3)
opg4j> var v4 = builder.addVertex("vertex 4").setProperty("double-
prop", 4.56789)

opg4j> builder.addEdge(v1, v2).setProperty("edge-prop",
"edge_prop_1_2").setLabel("label")
opg4j> builder.addEdge(v2, v3).setProperty("edge-prop",
"edge_prop_2_3").setLabel("label")
opg4j> builder.addEdge(v2, v4).setProperty("edge-prop",
"edge_prop_2_4").setLabel("label")
opg4j> builder.addEdge(v3, v4).setProperty("edge-prop",
"edge_prop_3_4").setLabel("label")
opg4j> builder.addEdge(v4, v2).setProperty("edge-prop",
"edge_prop_4_2").setLabel("label")

opg4j> var graph = builder.build()
```

#### Java

```
import oracle.pgx.api.*;
import oracle.pgx.common.types.IdType;

PgxSession session = Pgx.createSession("example");
GraphBuilder<String> builder =
session.createGraphBuilder(IdType.STRING);

VertexBuilder<String> v1 = builder.addVertex("vertex
```

```
1").setProperty("double-prop", 0.1);
VertexBuilder<String> v2 = builder.addVertex("vertex 2").setProperty("double-
prop", 2.0);
VertexBuilder<String> v3 = builder.addVertex("vertex 3").setProperty("double-
prop", 0.3);
VertexBuilder<String> v4 = builder.addVertex("vertex 4").setProperty("double-
prop", 4.56789);

builder.addEdge(v1, v2).setProperty("edge-prop",
"edge_prop_1_2").setLabel("label");
builder.addEdge(v2, v3).setProperty("edge-prop",
"edge_prop_2_3").setLabel("label");
builder.addEdge(v2, v4).setProperty("edge-prop",
"edge_prop_2_4").setLabel("label");
builder.addEdge(v3, v4).setProperty("edge-prop",
"edge_prop_3_4").setLabel("label");
builder.addEdge(v4, v2).setProperty("edge-prop",
"edge_prop_4_2").setLabel("label");

PgxGraph graph = builder.build();
```

## Python

```
from pypgx import get_session

session = get_session(session_name="example")
builder = session.create_graph_builder(id_type='string')

v1 = builder.add_vertex("vertex 1").set_property("double-prop", 0.1)
v2 = builder.add_vertex("vertex 2").set_property("double-prop", 2.0)
v3 = builder.add_vertex("vertex 3").set_property("double-prop", 0.3)
v4 = builder.add_vertex("vertex 4").set_property("double-prop", 4.56789)

builder.add_edge(v1, v2).set_property("edge-prop",
"edge_prop_1_2").set_label("label")
builder.add_edge(v2, v3).set_property("edge-prop",
"edge_prop_2_3").set_label("label")
builder.add_edge(v2, v4).set_property("edge-prop",
"edge_prop_2_4").set_label("label")
builder.add_edge(v3, v4).set_property("edge-prop",
"edge_prop_3_4").set_label("label")
builder.add_edge(v4, v2).set_property("edge-prop",
"edge_prop_4_2").set_label("label")

graph = builder.build()
```

---

### 14.6.1.6 Using Graph Builder with Implicit IDs

The `GraphBuilder` supports an `AUTO_GENERATED` generation strategy that allows to omit the edge or vertex IDs.

In this generation strategy, the graph server (PGX) will automatically assign IDs to the entities being added to the changeset. `PgxSession` supports `createGraphBuilder(IdGenerationStrategy vertexIdGenerationStrategy, IdGenerationStrategy edgeIdGenerationStrategy)` and `createGraphBuilder(IdType idType, IdGenerationStrategy vertexIdGenerationStrategy, IdGenerationStrategy edgeIdGenerationStrategy)` to specify the `IdGenerationStrategy`.

The following example illustrates creating a graph with three vertices and three edges using the `GraphBuilder` interface.

- 
- [JShell](#)
  - [Java](#)

## JShell

```
opg4j> var builder =
session.createGraphBuilder(IdGenerationStrategy.AUTO_GENERATED,
IdGenerationStrategy.AUTO_GENERATED)

opg4j> var v1 = builder.addVertex()
opg4j> var v2 = builder.addVertex()
opg4j> var v3 = builder.addVertex()

opg4j> builder.addEdge(v1, v2)
opg4j> builder.addEdge(v1, v3)
opg4j> builder.addEdge(v3, v2)

opg4j> var graph = builder.build()
```

## Java

```
import oracle.pgx.api.*;

PgxSession session = Pgx.createSession("example");
GraphBuilder<Integer> builder =
session.createGraphBuilder(IdGenerationStrategy.AUTO_GENERATED,
IdGenerationStrategy.AUTO_GENERATED);

VertexBuilder<Integer> v1 = builder.addVertex();
VertexBuilder<Integer> v2 = builder.addVertex();
VertexBuilder<Integer> v3 = builder.addVertex();
builder.addEdge(v1, v2);
builder.addEdge(v1, v3);
builder.addEdge(v3, v2);

PgxGraph graph = builder.build();
```

## 14.6.2 Modifying Loaded Graphs Using ChangeSet

This guide explains how to add and remove vertices and edges from already loaded graphs.

As a prerequisite, you must have a graph already loaded into the graph server (PGX). See [Reading Graphs from Oracle Database into the Graph Server \(PGX\)](#) for more information.

You can now use the `GraphChangeSet` interface to modify the loaded graphs.



### Note:

Modifying undirected graphs is not supported in graph server (PGX) 21.3.

- [Modifying Vertices](#)
- [Adding Edges](#)
- [GraphChangeSet with Partitioned IDs](#)
- [Error Handling when Using a ChangeSet](#)

### 14.6.2.1 Modifying Vertices

You can add, remove and modify vertices using the `GraphChangeSet` object.

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg4j> var changeSet = graph.<Integer>createChangeSet()

opg4j> changeSet.addVertex(42).setProperty("prop", 23)
opg4j> changeSet.updateVertex(128).setProperty("prop", 5)
opg4j> changeSet.removeVertex(1908)

opg4j> var updatedGraph = changeSet.build()

opg4j> updatedGraph.hasVertex(42) // Evaluates to: true
opg4j> updatedGraph.hasVertex(1908) // Evaluates to: false
```

#### Java

```
import oracle.pgx.api.*;

GraphChangeSet<Integer> changeSet = graph.createChangeSet();
```

```
changeSet.addVertex(42).setProperty("prop", 23);
changeSet.updateVertex(128).setProperty("prop", 5);
changeSet.removeVertex(1908);

PgxGraph updatedGraph = changeSet.build();
```

## Python

```
from pypgx.api import *

change_set = graph.create_change_set()

change_set.add_vertex(42).set_property("prop", 23)
changeSet.update_vertex(128).set_property("prop", 5)
changeSet.remove_vertex(1908)

updated_graph = change_set.build()
```

---

### 14.6.2.2 Adding Edges

You can also add edges to a graph using `GraphChangeSet`.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> var changeSet2 = updatedGraph.<Integer>createChangeSet()

opg4j> changeSet2.addEdge(333, 42).setProperty("cost", 42.3)
opg4j> changeSet2.addEdge(42, 99)

opg4j> var updatedGraph2 = changeSet2.build()
```

#### Java

```
import oracle.pgx.api.*;

GraphChangeSet<Integer> changeSet2 = graph.createChangeSet();

changeSet2.addEdge(333, 42).setProperty("cost", 42.42);
changeSet2.addEdge(42, 99);

PgxGraph updatedGraph2 = changeSet2.build();
```

## Python

```
from pypgx.api import *

change_set_2 = graph.create_change_set()
changeSet2.add_edge(333, 42).set_property("cost", 42.42)
changeSet2.add_edge(42, 99)
updated_graph_2 = change_set_2.build()
```

---

Note that by calling `changeSet2.build()`, you created a brand new graph with a unique name assigned by the graph server (PGX). If need be, you can specify a name argument to the `build()` method.

Additionally, you can create a new snapshot on top of the current graph with the `buildNewSnapshot()` method. See [Creating a Snapshot via ChangeSet](#) for more information.

### 14.6.2.3 GraphChangeSet with Partitioned IDs

You can use the `GraphChangeSet` API with graph with partitioned IDs. Ensure to set both the vertex ID generation strategy as well as the edge ID generation strategy to `IdGenerationStrategy.USER_IDS`. Furthermore, make sure to set the vertex ID type to `String`. An edge ID type does not need to be specified.

You can add, update and remove vertices and edges as shown in the following examples:

- 
- [Java](#)
  - [Python](#)

#### Java

```
GraphChangeSet<String> changeSet =
g.createChangeSet(IdGenerationStrategy.USER_IDS,
IdGenerationStrategy.USER_IDS);
changeSet.addVertex("Accounts(1002)").setProperty("NAME", "User1002");
changeSet.updateVertex("Accounts(4)").setProperty("NAME", "User4");
changeSet.removeVertex("Accounts(3)");
changeSet.addEdge("Transfers(5002)", "Accounts(5)",
"Accounts(6)").setProperty("AMOUNT", 12.50);
changeSet.updateEdge("Transfers(5)").setProperty("DESCRIPTION", 'Transfer
from User');
changeSet.removeEdge("Transfers(5001)");
PgxGraph g1 = changeSet.build();
```

#### Python

```
change_set = graph.create_change_set(vertex_id_generation_strategy =
'user_ids', edge_id_generation_strategy = 'user_ids')
change_set.add_vertex("Accounts(1002)").set_property("NAME", "User1002")
```

```
change_set.update_vertex("Accounts(4)").set_property("NAME", "User4")
change_set.remove_vertex("Accounts(3)")
change_set.remove_edge("Transfers(5001)")
PgxGraph g1 = change_set.build()
```

#### Note:

You cannot use the `setLabel()` API when IDs are partitioned. The vertex or edge will be labelled automatically based on the label attached to the provider (for which the name is provided as part of the ID). Similarly, you cannot set the vertex or edge key properties through the `setProperty()` API as the value is already extracted from the vertex or edge ID.

### 14.6.2.4 Error Handling when Using a ChangeSet

Error handling while populating a `ChangeSet` or while applying a `ChangeSet` to the existing graph can be configured by setting the `InvalidChangePolicy`. The options are:

- `OnInvalidChange.ERROR`: throws an exception (This is the default configuration)
- `OnInvalidChange.IGNORE`: ignores the issue and continues
- `OnInvalidChange.IGNORE_AND_LOG`: ignores the issue, logs in `DEBUG` log level and continues
- `OnInvalidChange.IGNORE_AND_LOG_ONCE`: only logs the first occurrence of each issue type

Issues that can be ignored with `InvalidChangePolicy` include trying to remove a vertex or an edge that does not exist in the graph, property type mismatch, updates to non existing properties, providing a vertex ID with wrong type or invalid vertex or edge providers.

The following example, tries to remove vertex 9032 which does not exist in the graph. By configuring `IGNORE_AND_LOG`, this action will be ignored while the property value update for vertex 99 will be applied successfully.

- [JShell](#)
- [Java](#)

#### JShell

```
opg4j> var changeSet3 = updatedGraph2.<Integer>createChangeSet()
opg4j>
changeSet3.setInvalidChangePolicy(OnInvalidChange.IGNORE_AND_LOG)

opg4j> changeSet3.removeVertex(9032)
opg4j> changeSet3.updateVertex(99).setProperty("prop1", 17)
```

```
opg4j> var updatedGraph3 = changeSet3.build() // will log that a vertex
removal was ignored

opg4j> var prop1Val = updatedGraph3.getVertex(99).getProperty("prop1") //
evaluates to 17
```

## Java

```
import oracle.pgx.api.*;

GraphChangeSet<Integer> changeSet3 = graph.createChangeSet();
changeSet3.setInvalidChangePolicy(OnInvalidChange.IGNORE_AND_LOG);
changeSet3.removeVertex(9032);
changeSet3.updateVertex(99).setProperty("prop1", 17);
PgxGraph updatedGraph3 = changeSet3.build(); // will log that a vertex
removal was ignored

int prop1Val = updatedGraph3.getVertex(99).getProperty("prop1"); //
evaluates to 17
```

### Note:

When connecting to a remote graph server (PGX), error handling log messages will not be relayed to the client. In such a case, you need access to the server logs to determine which issues have been ignored. For this, you must update the default log4j configuration file in `/etc/oracle/graph/log4j2.xml` and the graph server (PGX) logger configuration file in `/etc/oracle/graph/log4j2-server.xml` to log the DEBUG logs. You can then view the ignored issues in `/var/opt/log/pgx-server.log` file.

### Add Existing Edges and Vertices

The error handling for adding a vertex or an edge where its ID is already used in the graph or in an incompatible ChangeSet action can be configured with `AddExistingVertexPolicy` and `AddExistingEdgePolicy`.

### Note:

The default setting for `AddExistingVertexPolicy` and `AddExistingEdgePolicy` is `IGNORE`. This is different from `InvalidChangePolicy` where the default is `ERROR`.

## 14.7 Managing Transient Data

This guide discusses how to handle transient properties and collections.



The graph server (PGX) allows each client to maintain its own isolated workspace, called session. Clients may create additional data objects in their own session, which they can then use for analysis.

- [Managing Transient Properties](#)
- [Managing Collections and Scalars](#)

## 14.7.1 Managing Transient Properties

The graph server (PGX) adopts the Property Graph data model. Once a graph is loaded into PGX, the graph instance itself and its original properties are set as immutable. However, the client can create and attach additional properties to the graph dynamically. These extra properties are referred to as *transient* properties and are mutable by the client

The methods for creating transient properties are available in `PgxGraph`:

- 
- [Java](#)
  - [Python](#)

### Java

```
VertexProperty<ID, V> createVertexPropertyAsync(PropertyType type)
VertexProperty<ID, V> createVertexPropertyAsync(PropertyType type,
String name)
EdgeProperty<V> createEdgePropertyAsync(PropertyType type)
EdgeProperty<V> createEdgePropertyAsync(PropertyType type, String name)
```

In the preceding code:

- `PropertyType`: is an enum for the data type of the property, which must be one of the primitive types supported by PGX.
- `name`: is an optional argument to assign a unique name to the newly created property. If no name is specified, PGX will assign one to the client.

#### Note:

Names must be unique. There cannot be two different vertex or edge properties for the same graph and with the same name.

### Python

```
create_vertex_property(self, data_type, name=None)
```

---

All methods return a `Property` object, which represent the newly created transient property. Both of the underlying classes, `VertexProperty<ID, V>` and `EdgeProperty<V>`, are parametrized with the value type `V` the property holds. `V` matches the given `PropertyType`. `VertexProperty<ID, V>` is additionally parametrized with the vertex ID type. This is due to PGX support of several types of vertex identifiers. See our graph configuration chapter on how to specify the vertex ID type of a graph. `EdgeProperty<V>` is not parametrized with the edge ID type, because PGX only supports edge identifiers of type `long`.

- 
- [Java](#)
  - [Python](#)

## Java

```
GraphConfig config = GraphConfigBuilder.forFileFormats(...)
...
 .setVertexIdType(IdType.LONG)
...
 .build();

PgxGraph G = session.readGraphWithProperties(config);
VertexProperty<Long, String> p1 =
G.createVertexProperty(PropertyType.STRING);
EdgeProperty<Double> p2 = G.createEdgeProperty(PropertyType.DOUBLE);
```

## Python

```
G = session.read_graph_with_properties(config)
p1 = G.create_vertex_property("string")
p2 = G.create_edge_property("double")
```

---

To delete a transient property from the session, call `destroyAsync()` (or `destroy()`) on the property object.

## 14.7.2 Managing Collections and Scalars

The client can create graph-bound vertex and edge collections to use during the analysis with the following methods in `PgxGraph`:

- 
- [Java](#)
  - [Python](#)

## Java

```
VertexSequence<E> createVertexSequence()
VertexSequence<E> createVertexSequence(String name)
VertexSet<E> createVertexSet()
VertexSet<E> createVertexSet(String name)
EdgeSequence createEdgeSequence()
EdgeSequence createEdgeSequence(String name)
EdgeSet createEdgeSet()
EdgeSet createEdgeSet(String name)
```

## Python

```
create_edge_sequence(self, name=None)
create_vertex_sequence(self, name=None)
create_edge_set(self, name=None)
create_edge_sequence(self, name=None)
```

---

PGX also supports scalar collections such as `set` and `sequence`. Each of these collections can hold elements of various primitive data types like `INTEGER`, `LONG`, `FLOAT`, `DOUBLE` or `BOOLEAN`. Scalar collections are session-bound and can be created with the following methods in `PgxSession`:

```
ScalarSet<T> createSet(PropertyType contentType, String name)
ScalarSequence<T> createSequence(PropertyType contentType, String name)
ScalarSet<T> createSet(PropertyType contentType)
ScalarSequence<T> createSequence(PropertyType contentType)
```

In the preceding code, the optional argument (`name`) specifies the name of the newly created collection. If omitted, PGX chooses a name for the client. As with properties, the collections holding vertices are parametrized with the `ID` type of the vertices. Refer to graph configuration chapter to learn how to specify the vertex ID type of a graph.

The return value is the collection object which points to the newly created empty collection.

To drop a collection from the session, call `destroy()` on the collection object.

To check which collections are currently allocated for a graph you can use the following method:

- 
- [Java](#)
  - [Python](#)

## Java

```
Map<String, PgxCollection<? extends PgxEntity<?>, ?>> getCollections()
```

## Python

```
get_collections(self)
```

---

The returned map contains the names of the collections as keys and the collections as values. The collections can be casted to the matching collection subclass.

PGX supports special `Map` collection types and allows users to map between different data types (`oracle.pgx.common.types.PropertyType`). Maps can be created using `PgxGraph` or `PgxSession` APIs, the difference is that the latter supports only non graph-related types, and that the created maps directly depend on the session:

```
PgxMap<K, V> createMap(PropertyType keyType, PropertyType valType)
PgxMap<K, V> createMap(PropertyType keyType, PropertyType valType, String
mapName)
```

Similarly, scalar variables can be created in the client session using the following methods:

---

- [Java](#)
- [Python](#)

## Java

```
Scalar<T> createScalar(PropertyType type, String newScalarName)
Scalar<T> createScalar(PropertyType type)
```

## Python

```
create_scalar(self, data_type, name=None)
```

---

These collections and scalar variables can then be passed as arguments to graph algorithms. See [Using Custom PGX Graph Algorithms](#) for more information.

## 14.8 Graph Versioning

This guide describes the different ways to work with graph snapshots.

A graph can have multiple snapshots associated with it, reflecting different versions of the graph. All snapshots of a graph have the same graph configuration associated.

The following topics explain the various operations you can perform on graph snapshots:

- [Configuring the Snapshots Source](#)
- [Creating a Snapshot via Refreshing](#)

- [Creating a Snapshot via ChangeSet](#)
- [Checking Out the Latest Snapshots of a Graph](#)
- [Checking Out Different Snapshots of a Graph](#)
- [Directly Loading a Specific Snapshot of a Graph](#)

## 14.8.1 Configuring the Snapshots Source

Snapshots can be created from two sources: **Refreshing** and **ChangeSet**.

Refreshing is available for graphs that are read from a persistent data source, that is, a file. When the data source has changed with respect to the version stored in the graph server (PGX), it can be read again manually by calling the `PgxSession.readGraphWithProperties()` method. Similarly, if auto-refresh is set for the graph, the graph server (PGX) automatically reads the data source and creates new snapshots when the data source has changed.

Instead, a `ChangeSet` is a set of changes to a graph that the user creates and populates via the PGX `ChangeSet` API. Once a `ChangeSet` is created and populated with the desired changes, the user can simply call `GraphChangeSet.buildNewSnapshot()` to create a new snapshot for the graph. In this way, you are empowered to integrate changes coming from any source into the graph and build snapshots out of them.

Only one source of snapshots is allowed for a single graph and is chosen during graph configuration via the `snapshots_source` option, which can be set to either `REFRESH` or `CHANGE_SET`. In case the `snapshots_source` option is not explicitly set by the user, the following default settings apply:

- If the graph is from a persistent data source, the default value is `REFRESH`, so that snapshots can be created only by calling `PgxSession.readGraphWithProperties()` (or via auto-refresh, if configured).
- If the graph is transient, that is, built from a graph builder, the default value is `CHANGE_SET`, since the graph is not backed by a persistent data source from which changes can be read. It is for this reason, `CHANGE_SET` is the only admissible value for transient graphs.

Additionally, the following restrictions apply:

- If auto-refresh is enabled, then snapshots come from reading the backing data source and hence only `REFRESH` is admissible for the `snapshots_source` option.
- If the user attempts to create snapshots in a way that is different from the configuration (for example, by calling `GraphChangeSet.buildNewSnapshot()` when the graph's `snapshots_source` is `REFRESH`), the operation is invalid and an exception is thrown.

## 14.8.2 Creating a Snapshot via Refreshing

You can create a snapshot via refreshing by performing the following steps:

1. Create a session and load the graph into memory.
2. Check the available snapshots of the graph with `PgxSession.getAvailableSnapshots()` method.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> session.getAvailableSnapshots(G)
==> GraphMetaData [getNumVertices()=4, getNumEdges()=4, memoryMb=0,
dataSourceVersion=1453315103000, creationRequestTimestamp=1453315122669
(2016-01-20 10:38:42.669), creationTimestamp=1453315122685 (2016-01-20
10:38:42.685), vertexIdType=integer, edgeIdType=long]
```

## Java

```
Deque<GraphMetaData> snapshots = session.getAvailableSnapshots(G);
for(GraphMetaData metaData : snapshots) {
 System.out.println(metaData);
}
```

## Python

```
snapshots = session.get_available_snapshots(G)
for metadata in snapshots:
 print(metadata)
```

- 
3. Edit the source file to contain an additional vertex and an additional edge or insert two rows in the database.
  4. Reload the updated graph within the same session as you loaded the original graph. A new snapshot is created.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var G = session.readGraphWithProperties(G.getConfig(), true)
==> PGX Graph named 'sample_2' bound to PGX session 'a1744e86-65fb-4bd1-
b2dc-5458b20954a9' registered at PGX Server Instance running in embedded
mode
opg4j> session.getAvailableSnapshots(G)
==> GraphMetaData [getNumVertices()=4, getNumEdges()=4, memoryMb=0,
```

```
dataSourceVersion=1453315103000,
creationRequestTimestamp=1453315122669 (2016-01-20 10:38:42.669),
creationTimestamp=1453315122685 (2016-01-20 10:38:42.685),
vertexIdType=integer, edgeIdType=long]
==> GraphMetaData [getNumVertices()=5, getNumEdges()=5, memoryMb=3,
dataSourceVersion=1452083654000,
creationRequestTimestamp=1453314938744 (2016-01-20 10:35:38.744),
creationTimestamp=1453314938833 (2016-01-20 10:35:38.833),
vertexIdType=integer, edgeIdType=long]
```

## Java

```
G = session.readGraphWithProperties(G.getConfig(), true);

Deque<GraphMetaData> snapshots = session.getAvailableSnapshots(G);
```

## Python

```
G =
session.read_graph_with_properties(G.config,update_if_not_fresh=True
)
```

---

Note that there are two `GraphMetaData` objects in the call for available snapshots, one with 4 vertices and 4 edges and one with 5 vertices and 5 edges.

5. Verify that the graph variable points to the newly loaded graph using `getNumVertices()` and `getNumEdges()` methods.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> G.getNumVertices()
==> 5
opg4j> G.getNumEdges()
==> 5
```

## Java

```
int vertices = G.getNumVertices();
long edges = G.getNumEdges();
```

## Python

```
vertices = G.num_vertices
edges = G.num_edges
```

---

### 14.8.3 Creating a Snapshot via ChangeSet

You can create a graph snapshot with ChangeSet via the PGX Java API. When you want to create the graph from a persistent data source, you can use

`PgxSession.readGraphWithProperties()` with the `snapshots_source` configuration option set to `CHANGE_SET`.

You can create a snapshot via ChangeSet by performing the following steps:

1. Create a snapshot of a transient graph from database:
- 

- [JShell](#)
- [Java](#)
- [Python](#)

#### JShell

```
opg4j> var builder = session.createGraphBuilder()
opg4j> builder.addEdge(1, 2)
opg4j> builder.addEdge(2, 3)
opg4j> builder.addEdge(2, 4)
opg4j> builder.addEdge(3, 4)
opg4j> builder.addEdge(4, 2)
opg4j> var graph = builder.build()
```

#### Java

```
import oracle.pgx.api.*;

GraphBuilder<Integer> builder = session.createGraphBuilder();

builder.addEdge(1, 2);
builder.addEdge(2, 3);
builder.addEdge(2, 4);
builder.addEdge(3, 4);
builder.addEdge(4, 2);

PgxGraph graph = builder.build();
```



## Python

```
builder = session.create_graph_builder();

builder.add_edge(1, 2)
builder.add_edge(2, 3)
builder.add_edge(2, 4)
builder.add_edge(3, 4)
builder.add_edge(4, 2)

graph = builder.build()
```

- 
2. Create a `ChangeSet` from `graph` and populate it. The following example shows adding a new edge between vertices 1 and 4:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var changeSet = graph.<Integer>createChangeSet()
opg4j> changeSet.addEdge(6, 1, 4)
```

## Java

```
import oracle.pgx.api.*;
GraphChangeSet<Integer> changeSet = graph.createChangeSet();
changeSet.addEdge(6, 1, 4);
```

## Python

```
changeSet = graph.create_change_set() changeSet.add_edge(1,4,6)
```

- 
3. Create a second snapshot using `GraphChangeSet.buildNewSnapshot()` as shown in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var secondSnapshot = changeSet.buildNewSnapshot()
opg4j> session.getAvailableSnapshots(secondSnapshot).size()
==> 2
```

## Java

```
PgxGraph secondSnapshot = changeSet.buildNewSnapshot();
System.out.println(session.getAvailableSnapshots(secondSnapshot).size())
;
```

## Python

```
second_snapshot = change_set.build_new_snapshot()
print(len(session.get_available_snapshots()))
```

---

Thus two snapshots, referenced via the variables `graph` and `secondSnapshot` are created.

## 14.8.4 Checking Out the Latest Snapshots of a Graph

With multiple snapshots of a graph being available and regardless of their source, you can check out a specific snapshot using the `PgxSession.setSnapshot()` method. You can use the `LATEST_SNAPSHOT` constant of `PgxSession` to easily check out the latest available snapshot, as shown in the following example:

- 
- [JShell](#)
  - [Java](#)

## JShell

```
opg4j> session.setSnapshot(G, PgxSession.LATEST_SNAPSHOT)
==> null
opg4j> session.getCreationTimestamp()
==> 1453315122685
```

## Java

```
session.setSnapshot(G, PgxSession.LATEST_SNAPSHOT);
System.out.println(session.getCreationTimestamp());
```

---

See the printed timestamp to verify the most recent snapshot.

## 14.8.5 Checking Out Different Snapshots of a Graph

You can also check out a specific snapshot, again using the `PgxSession.setSnapshot()`.

For example, consider the following two snapshots of a graph:

```
==> GraphMetaData [getNumVertices()=4, getNumEdges()=4, memoryMb=0,
dataSourceVersion=1453315103000,
creationRequestTimestamp=1453315122669 (2016-01-20 10:38:42.669),
creationTimestamp=1453315122685 (2016-01-20 10:38:42.685),
vertexIdType=integer, edgeIdType=long]
==> GraphMetaData [getNumVertices()=5, getNumEdges()=5, memoryMb=3,
dataSourceVersion=1452083654000,
creationRequestTimestamp=1453314938744 (2016-01-20 10:35:38.744),
creationTimestamp=1453314938833 (2016-01-20 10:35:38.833),
vertexIdType=integer, edgeIdType=long]
```

To check out a specific snapshot of the graph, you must pass the `creationTimestamp` of the snapshot you want to load to `setSnapshot()`.

For example, if `G` is pointing to the newest graph with 5 vertices and 5 edges, but you want to analyze the older graph, you need to set the snapshot to 1453315122685.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
opg4j> G.getNumVertices()
==> 5
opg4j> G.getNumEdges()
==> 5
opg4j> session.setSnapshot(G, 1453315122685)
==> null
opg4j> G.getNumVertices()
==> 4
opg4j> G.getNumEdges()
==> 4
```

### Java

```
session.setSnapshot(G,1453315122685);
```

## Python

```
session.set_snapshot(G,1453315122685)
```

---

Note that setting the snapshot, changes the number of vertices and edges from 5 to 4.

Alternatively, you can also retrieve the creation timestamp of each snapshot from its associated `GraphMetaData` object via the `GraphMetaData.getCreationTimestamp()` method. The easiest way to get the `GraphMetaData` information of all the snapshots is to use the `PgxSession.getAvailableSnapshots()` method, which returns a collection of `GraphMetaData` information of each snapshot ordered by creation timestamp from the most recent to the oldest.

### 14.8.6 Directly Loading a Specific Snapshot of a Graph

You can also load a specific snapshot of a graph directly using the `PgxSession.readGraphAsOf()` method. This is a shortcut for loading a graph with `readGraphWithProperties()` followed by a `setSnapshot()`.

Consider two snapshots of a graph that are already loaded into the PGX session. The following example shows how to get a reference to a specific snapshot:

1. Get a graph configuration for the graph:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> var config =
GraphConfigFactory.forAnyFormat().fromPath("<path_to_json>")
==> {"format":"adj_list", ... }
```

#### Java

```
GraphConfig config =
GraphConfigFactory.forAnyFormat().fromPath("<path_to_json>");
```

#### Python

```
config = GraphConfigFactory.for_any_format().from_path("<path_to_json>")
```

---

2. Check the loaded snapshots for this graph config using `getAvailableSnapshots()`:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> session.getAvailableSnapshots(G)
==> GraphMetaData [getNumVertices()=4, getNumEdges()=4, memoryMb=0,
dataSourceVersion=1453315103000,
creationRequestTimestamp=1453315122669 (2016-01-20 10:38:42.669),
creationTimestamp=1453315122685 (2016-01-20 10:38:42.685),
vertexIdType=integer, edgeIdType=long]
==> GraphMetaData [getNumVertices()=5, getNumEdges()=5, memoryMb=3,
dataSourceVersion=1452083654000,
creationRequestTimestamp=1453314938744 (2016-01-20 10:35:38.744),
creationTimestamp=1453314938833 (2016-01-20 10:35:38.833),
vertexIdType=integer, edgeIdType=long]
```

## Java

```
Deque<GraphMetaData> snapshots = session.getAvailableSnapshots(G);
```

## Python

```
session.get_available_snapshots(G)
```

3. Check out the snapshot of the graph which has 4 vertices and 4 edges and having the timestamp 1453315122685:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var G = session.readGraphAsOf(config, 1453315122685)
==> PGX Graph named 'sample' bound to PGX session
'a1744e86-65fb-4bd1-b2dc-5458b20954a9' registered at PGX Server
Instance running in embedded mode
opg4j> G.getNumVertices()
==> 4
```

```
opg4j> G.getNumEdges()
==> 4
```

## Java

```
PgxGraph G = session.readGraphAsOf(config, 1453315122685);
```

## Python

```
G = read_graph_as_of(config, creation_timestamp=1453315122685)
```

---

## 14.9 Labels and Properties

You can perform various actions on the graph property and label values by executing PGQL queries.

- [Setting and Getting Property Values](#)
- [Getting Label Values](#)

### 14.9.1 Setting and Getting Property Values

#### Getting Property Values

You can obtain the vertex or edge property values by executing a `SELECT` PGQL query on the graph.

For example:

- 
- [JShell](#)
  - [Java](#)

#### JShell

```
opg4j> session.queryPgql("SELECT e.src_id, e.dest_id, e.amount FROM MATCH
(n:Account) -[e:Transfers]-> (m:Account) on bank_graph").print()
```

#### Java

```
...
...
PgxGraph g = session.getGraph("bank_graph");
String query =
 "SELECT e.src_id, e.dest_id, e.amount FROM MATCH (n:Account) -
[e:Transfers]-> (m:Account)";
g.queryPgql(query).print();
```

The resulting property values may appear as:

```
+-----+
| src_id | dest_id | amount |
+-----+
| 1 | 259 | 1000 |
| 1 | 418 | 1000 |
| 1 | 584 | 1000 |
| 1 | 644 | 1000 |
| 1 | 672 | 1000 |
| 2 | 493 | 1000 |
| 2 | 546 | 1000 |
| 2 | 693 | 1000 |
| 2 | 833 | 1000 |
| 2 | 840 | 1000 |
+-----+
```

### Setting Property Values

You can set the vertex or edge property values by executing insert or update PGQL queries on the graph.

For example, to set a new vertex account ID on a graph using `INSERT` query:

- [JShell](#)
- [Java](#)

### JShell

```
opg4j> PgxGraph g = session.getGraph("bank_graph_analytics")
g ==>
PgxGraph[name=bank_graph_analytics,N=1000,E=5001,created=1616312153556]
opg4j> PgxGraph g_mutable = g.clone("bank_graph_analytics_copy")
g_mutable ==>
PgxGraph[name=bank_graph_analytics_copy,N=1000,E=5001,created=161631241
3799]
opg4j> g_mutable.executePgql("INSERT VERTEX v LABELS (Accounts)
PROPERTIES (v.id = 1001)")
```

### Java

```
...
...
PgxGraph g1 =
session.readGraphWithProperties("bank_graph_analytics.json");
PgxGraph g2 = g1.clone("bank_graph_analytics_copy");
g2.executePgql("INSERT VERTEX v " +
```

```
" LABELS (Accounts) " +
" PROPERTIES (v.id = 1001)");
```

## 14.9.2 Getting Label Values

You can retrieve the vertex or edge label values of a graph as shown:

```
PgxGraph g = session.getGraph("bank_graph_analytics");
String query =
 "SELECT LABEL(v), COUNT(*) "
 + "FROM MATCH (v) "
 + "GROUP BY LABEL(v) "
 + "ORDER BY COUNT(v) DESC";
PgqlResultSet resultSet = g.queryPgql(query);
resultSet.print();
```

The result may appear as shown:

```
+-----+
| LABEL(n) | COUNT(*) |
+-----+
| ACCOUNT | 1000 |
+-----+
```

## 14.10 Filter Expressions

This guide explains the usage of filter expressions.

Filter expressions are applied in the following scenarios:

- **Path-Finding:** Include only specific vertices and edges in a path
- **Sub-Graphs:** Include only specific vertices and edges in a subgraph
- **Set creation:** Create a vertex or edge set and include only specific vertices or edges

There are two types of filter expressions:

- **Vertex filters:** Evaluated on each vertex
- **Edge filters:** Evaluated on each edge, including the two vertices it connects.

These filter expressions will evaluate to `true` if the current edge or vertex matches the expression or to `false` if it does not. Filter expressions are stateless and side-effect free.

The following short example below will evaluate to `true` for all edges where the source vertex's string property name is "PGX".

```
src.name="PGX"
```

- [Syntax](#)
- [Type System](#)



- [Path Finding Filters](#)
- [Subgraph Filters](#)
- [Operations on Filter Expressions](#)

## 14.10.1 Syntax

### Trivial Expressions

Always evaluates to `true`:

```
true
```

Always evaluates to `false`:

```
false
```

### Constants

Legal constants are integer, long and floating point numbers of single and double precision as well as strings literals and `true` and `false`. Long constants need to be suffixed with `L` or `L`. Floating point numbers are treated as double precision numbers by default. To force a certain precision you can use `f` or `F` for single precision and `d` or `D` for double precision floating point numbers. String literals are UTF-8 character sequences, surrounded by single or double quotation marks.

```
25
4294967296L
0.62f
0.33d
"Double quoted string"
'Single quoted string'
```

### Vertex and Edge Identifiers

Depending on the filter type, different identifiers are valid.

#### Vertex Filter

Vertex filter expressions have only one keyword that addresses the vertex in the current context.

`vertex` denotes the vertex that is currently being evaluated by the filter expression.

```
vertex
```

#### Edge Filter

Edge filter expressions have several keywords that addresses the edge or its vertices in the current context.

`edge` denotes the edge that is currently being evaluated by the filter expression.

```
edge
```

`dst` denotes the destination vertex of the current edge. `dst` is only valid in the subgraph context.

`dst`

`src` denotes the source vertex of the current edge. `src` is only valid in the subgraph context.

`src`

## Properties

Filter expressions can access the values of vertex and edge properties.

`<id>.<property>`

where:

- `<id>`: is any vertex or edge identifier (that is, `src`, `dst`, `vertex`, `edge`).
- `<property>`: is the name of a vertex or edge property.

### Note:

This has to be the name of an edge property if the identifier is `edge`. Otherwise it has to be a vertex property.

If the property name is a reserved name in the filter expression syntax or contains spaces, it must be quoted in single or double quotes.

The following code accesses the 'cost' property of the source vertex.

`src.cost`

Temporal properties support values comparison (constants and property values) using special constructors. The default temporal formats are shown in the following table:

**Table 14-5 Default Temporal Formats**

Property Type	Constructor
DATE	<code>date ('yyyy-MM-dd HH:mm:ss')</code>
LOCAL_DATE	<code>date 'yyyy-MM-dd'</code>
TIME	<code>time 'HH:mm:ss'</code>
TIME_WITH_TIMEZONE	<code>time 'HH:mm:ss+/-XXX'</code>
TIMESTAMP	<code>timestamp 'yyyy-MM-dd HH:mm:ss'</code>
TIMESTAMP_WITH_TIMEZONE	<code>timestamp 'yyyy-MM-dd HH:mm:ss+/-XXX'</code>

The following expression accesses the property 'timestamp\_withTZ' of an edge and checks if it is equal to 3/27/2007 06:00+01:00.

`edge.timestamp_withTZ = timestamp'2007-03-2706:00:00+01:00'`

 **Note:**

*Properties* of type *date* can only be checked for equality. *date* type usage is deprecated since version 2.5, instead use *local date* or *timestamp* types that support all operations.

**Methods**

Filter expressions support the following functions:

**Degree Functions**

1. `outDegree()` returns the number of outgoing edges of the vertex identifier. `degree()` is a synonym for `outDegree`.

```
int <id>.degree()
int <id>.outDegree()
```

The following example determines whether the `out-degree` of the source vertex is greater than three:

```
src.degree() > 3
```

2. `inDegree()` returns the number of incoming edges of the vertex identifier.

```
int <id>.inDegree()
```

**Label Functions**

1. `hasLabel()` checks if a vertex has a label.

```
boolean <id>.hasLabel('<label>')
```

The following example determines whether a vertex has the label "city":

```
vertex.hasLabel('city')
```

2. `label()` returns the label of an edge.

```
string <id>.label()
```

The following expression checks whether the label of an edge is "clicked\_by":

```
edge.label() = 'clicked_by'
```

**Relational Expressions**

To compare values (e.g., property values or constants), filter expressions provide the comparison operators listed below. Note: Both `==` and `=` are synonyms.

```
==
=
!=
<
<=
>
>=
```

The following example checks whether the "cost" property of the source vertex is lower than or equals to 1.23.

```
src.cost <= 1.23
```

## Vertex ID Comparison

It is also possible to filter for vertices with a specific vertex ID.

```
<id> = <vertex_id>
```

The following example determines whether the source vertex of an edge has the vertex ID "San Francisco"

```
src = "San Francisco"
```

## Regular Expressions

Strings can be matched using regular expressions.

```
<string expression> =~ '<regexexpression>'
```

The following example checks if the edge label starts with a lowercase letter and ends with a number:

```
edge.label() =~ '^[a-z].*[0-9]$'
```



### Note:

The syntax followed for the pattern on the right-hand side, is [Java REGEX](#).

## Type Conversions

The following syntax allows converting the type of <expression> to <type>.

```
(<type>) <expression>
```

The following example converts the value of the 'cost' property of the source vertex to an integer value:

```
(int) src.cost
```

## Boolean Expressions

Filter expressions can be composed to form other filter expressions. This can be done using the Boolean operators && (and), || (or) and ! (not).



### Note:

Only boolean operands can be composed.

```
(! true) || false
edge.cost < INF && dst.visited = false
src.degree() < 10 || !(dst.visited)
```

## Arithmetic Expressions

Any numeric expression can be combined using arithmetic expressions. The available arithmetic operators are: +, -, \*, /, %.

**Note:**

These operators only work on numeric operands.

```
1 + 5
-vertex.degree()
edge.cost * 2 > 5
src.value * 2.5 = (dst.inDegree() + 5) / dst.outDegree()
```

**Operator Precedence**

Operator precedences are shown in the following list, from highest precedence to the lowest. An operator on a higher level is evaluated before an operator on a lower level.

1. + (unary plus), - (unary minus)
2. \*, /, %
3. +, -
4. =, !=, <, >, <=, >=, =~
5. NOT
6. AND
7. OR

**Syntactic Sugar**

`both` and `any` denote the source and destination vertex of the current edge. They can be used to express a condition that should be `true` for both or at least either one of the two vertices. These keywords are only valid in an edge filter expression. To use them in a vertex filter results in a runtime type-checking exception.

```
both
any
```

The filter expressions inside the following examples are equivalent:

```
both.property = 1
src.property = 1 && dst.property = 1

any.degree() > 1
src.degree() > 1 || dst.degree() > 1
```

## 14.10.2 Type System

Filter expressions are a very simple type system. There are only the following 13 types:

1. `integer` (can be abbreviated in expressions with `int`)
2. `long`
3. `float`
4. `double`
5. `boolean`

6. string
7. date
8. time
9. time with timezone
10. timestamp
11. timestamp with timezone
12. vertex
13. edge

Conversions are only allowed from one numeric type to another numeric type (i.e. integer, float, double, long).

Comparisons require both sides to be of the same (or convertible) type.

### 14.10.3 Path Finding Filters

Filters can be used to limit the analyzed edges when searching for a shortest path between a source and destination vertex in a graph.

An edge filter expression is evaluated against each edge that is visited during the traversal of the graph. If the filter evaluates to `false` on an edge, this edge will be ignored and will not appear in the resulting shortest path.

It is also possible to use a vertex filter for path finding.

A vertex filter expression is evaluated against each vertex that is visited during the traversal of the graph, except for the source and destination vertex.

If the filter evaluates to `false` on a vertex, the edge to this vertex and all outgoing edges of the vertex will be ignored. The vertex will not appear in the resulting shortest path.

The source and destination vertex can be any vertex in the graph and the filter is not evaluated for them.

### 14.10.4 Subgraph Filters

#### Edge Filters

An edge filter expression is evaluated for each edge in the graph. The edge filter has access to the source and destination vertex of each edge and all of its properties.

If the filter expression evaluates to true, the edge and both the source and destination vertex will appear in the subgraph.

#### Vertex Filters

A vertex filter expression is evaluated for every vertex in the graph.

Every vertex for which the filter expression evaluates to true will appear in the subgraph.

Every edge connecting two vertices for which the expression evaluates to true will also appear in the subgraph.

#### Result Set Filters

Result set edge and vertex filters allow the creation of edge and vertex sets out of a given PGQL result set.

### Vertex and Edge Collection Filters

Vertex and edge collection filters allow the creation of edge and vertex filters out of a given vertex and edge collection.

## 14.10.5 Operations on Filter Expressions

This section explains the various operations that you can perform on filter expressions.

- [Defining Filter Expressions](#)
- [Defining Result Set Filters](#)
- [Creating a Subgraph from PGQL Result Set](#)
- [Defining Collection Filters](#)
- [Creating a Subgraph from Collection Filters](#)
- [Combining Filter Expressions](#)
- [Creating a Subgraph Using Filter Expressions with Partitioned IDs](#)

### 14.10.5.1 Defining Filter Expressions

You can define a new vertex filter, as shown in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> var vertexFilter = VertexFilter.fromExpression("vertex.name = 'PGX'")
```

#### Java

```
VertexFilter vertexFilter = VertexFilter.fromExpression("vertex.name = 'PGX'");
```

#### Python

```
vertex_filter = VertexFilter("vertex.name = 'PGX'")
```

---

You can define a new edge filter, as shown in the following code:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var edgeFilter = EdgeFilter.fromExpression("edge.cost > 5")
```

## Java

```
EdgeFilter edgeFilter = EdgeFilter.fromExpression("edge.cost > 5");
```

## Python

```
vertex_filter = EdgeFilter("edge.cost > 5")
```

---

### 14.10.5.2 Defining Result Set Filters

You can define a result set vertex filter, as shown in the following code:

- 
- [JShell](#)
  - [Java](#)

## JShell

```
// Evaluates query on graph g to obtain a result set
opg4j> var resultSet = g.queryPgql("SELECT x FROM MATCH (x) WHERE x.age >
24")
// Define a filter on the result set for the column "x"
opg4j> var vertexFilter = VertexFilter.fromPgqlResultSet(resultSet, "x")
// Obtain a vertex set
opg4j> var vertexSet = g.getVertices(vertexFilter)
```

## Java

```
// Evaluates query on graph g to obtain result set
PgqlResultSet resultSet = g.queryPgql("SELECT x FROM MATCH (x) WHERE x.age >
24");
// Define a filter on the result set for the column "x"
VertexFilter vertexFilter = VertexFilter.fromPgqlResultSet(resultSet, "x");
```



```
// Obtain a vertex set
VertexSet vertexSet = g.getVertices(vertexFilter);
```

---

You can define a result set edge filter, as shown in the following code:

---

- [JShell](#)
- [Java](#)

## JShell

```
// Evaluates query on graph g to obtain result set
opg4j> var resultSet = g.queryPgql("SELECT e FROM MATCH ()-[e]->()
WHERE e.weight >= 8")
// Define a filter on the result set for the column "e"
opg4j> var edgeFilter = EdgeFilter.fromPgqlResultSet(resultSet, "e")
// Obtain an edge set
opg4j> var edgeSet = g.getEdges(edgeFilter)
```

## Java

```
// Evaluates query on graph g to obtain result set
PgqlResultSet resultSet = g.queryPgql("SELECT e FROM MATCH ()-[e]->()
WHERE e.weight >= 8");
// Define a filter on the result set for the column "e"
EdgeFilter edgeFilter = EdgeFilter.fromPgqlResultSet(resultSet, "e");
// Obtain an edge set
EdgeSet edgeSet = g.getEdges(edgeFilter);
```

---

### 14.10.5.3 Creating a Subgraph from PGQL Result Set

A subgraph can be obtained from a PGQL result set using result set filters.

You can create a subgraph from a result set vertex filter, as shown in the following code:

---

- [JShell](#)
- [Java](#)

## JShell

```
// Evaluates query on graph g to obtain result set
opg4j> var resultSet = g.queryPgql("SELECT x FROM MATCH (x) WHERE x.age >
24")
// Define a filter on the result set for the column "x"
opg4j> var resultSetVertexFilter = VertexFilter.fromPgqlResultSet(resultSet,
"x")
// Create a subgraph of g containing the matched vertices in the resultSet
and the edges that connect them if any.
opg4j> var newGraph = g.filter(resultSetVertexFilter)
```

## Java

```
// Evaluates query on graph g to obtain result set
PgqlResultSet resultSet = g.queryPgql("SELECT x MATCH (x) WHERE x.age > 24");
// Define a filter on the result set for the column "x"
VertexFilter resultSetVertexFilter =
VertexFilter.fromPgqlResultSet(resultSet, "x");
// Create a subgraph of g containing the matched vertices in the resultSet
and the edges that connect them if any.
PgxGraph newGraph = g.filter(resultSetVertexFilter);
```

---

You can create a subgraph from a result set edge filter, as shown in the following code:

---

- [JShell](#)
- [Java](#)

## JShell

```
// Evaluates query on graph g to obtain result set
opg4j> var resultSet = g.queryPgql("SELECT e FROM MATCH ()-[e]->() WHERE
e.cost < 100")
// Define a filter on the result set for the column "e"
opg4j> var resultSetEdgeFilter = EdgeFilter.fromPgqlResultSet(resultSet, "e")
// Create a subgraph of g containing the matched edges in the resultSet and
their corresponding source and destination vertices.
opg4j> var newGraph = g.filter(resultSetEdgeFilter)
```

## Java

```
// Evaluates query on graph g to obtain result set
PgqlResultSet resultSet = g.queryPgql("SELECT e FROM MATCH ()-[e]->() WHERE
e.cost < 100");
// Define a filter on the result set for the column "e"
EdgeFilter resultSetEdgeFilter = EdgeFilter.fromPgqlResultSet(resultSet,
"e");
```

```
// Create a subgraph of g containing the matched edges in the
resultSet and their corresponding source and destination vertices.
PgxGraph newGraph = g.filter(resultSetEdgeFilter);
```

---

## 14.10.5.4 Defining Collection Filters

You can define a vertex collection filter, as shown in the following code:

- [JShell](#)
- [Java](#)

### JShell

```
// Obtain a vertex collection from an algorithm, query execution or
any other way
opg4j> VertexCollection<?> vertexCollection = ...
// Define a filter from the collection
opg4j> var vertexFilter = VertexFilter.fromCollection(vertexCollection)
```

### Java

```
// Obtain a vertex collection from an algorithm, query execution or
any other way
VertexCollection<?> vertexCollection = ...
// Define a filter from the collection
VertexFilter vertexFilter =
VertexFilter.fromCollection(vertexCollection);
```

---

You can define an edge collection filter, as shown in the following code:

- [JShell](#)
- [Java](#)

### JShell

```
// Obtain an edge collection from an algorithm, query execution or any
other way
opg4j> EdgeCollection edgeCollection = ...
// Define a filter from the collection
opg4j> var edgeFilter = EdgeFilter.fromCollection(edgeCollection)
```

## Java

```
// Obtain an edge collection from an algorithm, query execution or any other way
EdgeCollection edgeCollection = ...
// Define a filter from the collection
EdgeFilter edgeFilter = EdgeFilter.fromCollection(edgeCollection);
```

---

### 14.10.5.5 Creating a Subgraph from Collection Filters

A subgraph can be obtained by using vertex or edge collection filters.

You can create a subgraph from vertex collection filter, as shown in the following code:

---

- [JShell](#)
- [Java](#)

#### JShell

```
// Obtain a vertex collection from an algorithm, query execution or any other way
opg4j> VertexCollection<?> vertexCollection = ...
// Define a filter from the collection
opg4j> var vertexFilter = VertexFilter.fromCollection(vertexCollection)
// Create a subgraph of g containing the matched vertices in the vertex collection and the edges that connect them if any.
opg4j> var newGraph = g.filter(vertexFilter)
```

#### Java

```
// Obtain a vertex collection from an algorithm, query execution or any other way
VertexCollection<?> vertexCollection = ...
// Define a filter from the collection
VertexFilter vertexFilter = VertexFilter.fromCollection(vertexCollection);
// Create a subgraph of g containing the matched vertices in the vertex collection and the edges that connect them if any.
PgxGraph newGraph = g.filter(vertexFilter);
```

---

You can create a subgraph from edge collection filter, as shown in the following code:

---

- [JShell](#)

- Java

## JShell

```
// Obtain an edge collection from an algorithm, query execution or any
other way
opg4j> EdgeCollection edgeCollection = ...
// Define a filter from the collection
opg4j> var edgeFilter = EdgeFilter.fromCollection(edgeCollection)
// Create a subgraph of g containing the matched edges in the
collection and their corresponding source and destination vertices.
opg4j> var newGraph = g.filter(edgeFilter)
```

## Java

```
// Obtain an edge collection from an algorithm, query execution or any
other way
EdgeCollection edgeCollection = ...
// Define a filter from the collection
EdgeFilter edgeFilter = EdgeFilter.fromCollection(edgeCollection);
// Create a subgraph of g containing the matched edges in the
collection and their corresponding source and destination vertices.
PgxGraph newGraph = g.filter(edgeFilter);
```

---

### 14.10.5.6 Combining Filter Expressions

Any filter expression used for subgraph filtering, can be combined with any other filter expression to form a new filter expression.

Filters can be combined using the following operations:

- intersection
- union

The intersection of two filters will only keep a vertex or edge, if both filters would accept it.



#### Note:

The intersection of two filters will not behave as an **AND** in the filter expression.

The union of two filters will keep a vertex or edge, if one of the filters would accept it.



#### Note:

The union of filters will not behave as an **OR** in the filter expression.

In the following example an edge filter is intersected with a vertex filter. The resulting subgraph will only include vertices that have the name 'PGX' and will only include edges that have a cost greater than 5.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> var edgeFilter = EdgeFilter.fromExpression("edge.cost > 5")
opg4j> var vertexFilter = VertexFilter.fromExpression("vertex.name = 'PGX'")
opg4j> var combinedFilter = edgeFilter.intersect(vertexFilter)
```

## Java

```
EdgeFilter edgeFilter = EdgeFilter.fromExpression("edge.cost > 5");
VertexFilter vertexFilter = VertexFilter.fromExpression("vertex.name =
'PGX'");
GraphFilter combinedFilter = edgeFilter.intersect(vertexFilter);
```

## Python

```
edge_filter = EdgeFilter("edge.cost > 5")
vertex_filter = VertexFilter("vertex.name = 'PGX'")
combined_filter = edge_filter.intersect(vertex_filter)
```

---

In contrast, the subgraph created by the union of those filters will include vertices that either have the name 'PGX' or that has an incoming or outgoing edge with a cost greater than 5. It will also include edges with a cost greater than 5, as well as edges for which the source and destination vertex have the name 'PGX'.

### 14.10.5.7 Creating a Subgraph Using Filter Expressions with Partitioned IDs

You can create a subgraph using filter expressions with partitioned IDs.

For example, the following creates a subgraph that contains only a single vertex with ID Account(1):

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
opg4j> PgxGraph subgraph =
g.filter(VertexFilter.fromExpression("vertex = 'Accounts(1) '"))
subgraph ==> PgxGraph[name=sub-graph_26,N=1,E=0,created=1630414040396]
```

## Java

```
PgxGraph subgraph = g.filter(VertexFilter.fromExpression("vertex =
'Accounts(1) '"));
```

## Python

```
subgraph = graph.filter(VertexFilter.from_expression("vertex =
'Accounts(1) '"))
```

---

The following example creates a subgraph that contains only a single edge with ID Transfers(1), and two accompanying vertices:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> PgxGraph subgraph = g.filter(EdgeFilter.fromExpression("edge =
'Transfers(1) '"))
subgraph ==> PgxGraph[name=sub-graph_27,N=2,E=1,created=1630414144529]
```

## Java

```
PgxGraph subgraph = g.filter(EdgeFilter.fromExpression("edge =
'Transfers(1) '"));
```

## Python

```
subgraph = graph.filter(EdgeFilter.from_expression("edge =
'Transfers(1) '"))
```

## 14.11 Advanced Task Scheduling Using Execution Environments

This guide shows how you can use the advanced scheduling features of the enterprise scheduler.

The enterprise scheduler features of the graph server (PGX) are currently only available for Linux (x86\_64), macOS (x86\_64) and Solaris (x86\_64, sparc).

The following topics provide more detailed information on enabling and scheduling tasks using the execution environment:

- [Enterprise Scheduler Configuration Guide](#)
- [Enabling Enterprise Scheduler Features](#)
- [Retrieving and Inspecting the Execution Environment](#)
- [Modifying and Submitting Tasks Under an Updated Environment](#)
- [Using Lambda Syntax](#)

### 14.11.1 Enterprise Scheduler Configuration Guide

This chapter describes the extra configuration options for the enterprise scheduler.



#### Note:

These configuration options are only available if the `scheduler` configuration variable is set to `enterprise_scheduler` in [Configuration Parameters for the Graph Server \(PGX\) Engine](#).

The configuration is divided into the following two parts:

1. `enterprise_scheduler_config`: for setting details about how tasks should be scheduled
2. `enterprise_scheduler_flags`: where you can configure the enterprise scheduler in more detail

#### Enterprise Scheduler Fields

Field	Type	Description	Default
<code>analysis_tas_k_config</code>	<code>object</code>	Configuration for analysis tasks.	<b>weight</b> <no-of-CPU>  <b>priority</b> medium  <b>max_threads</b> <no-of-CPU>



Field	Type	Description	Default
fast_analysis_task_config	object	Configuration for fast analysis tasks.	<b>weight</b> 1  <b>priority</b> high  <b>max_threads</b> <no-of-CPUs>
max_num_concurrent_io_tasks	integer	Maximum number of concurrent io tasks at a time.	3
num_io_threads_per_task	integer	Number of io threads to use per task.	<no-of-cpus>

### Analysis Task Config Fields

Field	Type	Description	Default
max_threads	integer	A hard limit on the number of threads to use for a task.	<i>required</i>
priority	enum[high, medium, low]	The priority of the task. Threads are given to the task with the highest priority at the moment of execution. If there are more threads that have the highest priority, threads are given to the tasks according to their weight	<i>required</i>
weight	integer	The weight of the task. Threads are given to tasks proportionally to their weight. Tasks with higher weight will get more threads than tasks with lower weight. Tasks with the same weight will get the same amount of threads.	<i>required</i>

### Enterprise Scheduler Flags

Field	Type	Description	Default
show_allocations	boolean	If <code>true</code> show memory allocation information.	false
show_environment	boolean	If <code>true</code> show version numbers and main environment settings at startup.	false
show_logging	boolean	If <code>true</code> enable summary logging. This is available even in non-debug builds and includes information such as the machine hardware information obtained at start-up, and per-job / per-loop information about the workload.	false
show_profiling	boolean	If <code>true</code> show profiling information.	false
show_scheduler_state	boolean	If <code>true</code> dump scheduler state on each update.	false

Field	Type	Description	Default
show_warnings	boolean	If true enable warnings. These are non-fatal errors. For example, if a NUMA-aware allocation cannot be placed on the intended socket.	true

#### Example 14-4 Custom Enterprise Scheduler Configuration

This configuration sets the number of io threads per task to 16, increases the maximum number of concurrent io tasks to 5. It also sets the configuration for fast analysis tasks to have a weight of 1, priority of "high" and sets a limit to the maximum number of threads used to 1.

```
{
 "enterprise_scheduler_config": {
 "num_io_threads_per_task": 16,
 "max_num_concurrent_io_tasks": 5,
 "fast_analysis_task_config": {
 "weight": 1,
 "priority": "high",
 "max_threads": 1
 }
 }
}
```

#### Example 14-5 Using the Enterprise Scheduler Flags

This configuration enables extra logging output from the enterprise scheduler.

```
{
 "enterprise_scheduler_flags": {
 "show_logging": true
 }
}
```

## 14.11.2 Enabling Enterprise Scheduler Features

You can enable the enterprise scheduler features, by setting the flag `allow_override_scheduling_information` of the the graph server (PGX) configuration file to true:

```
{"allow_override_scheduling_information":true}
```

See [Configuration Parameters for the Graph Server \(PGX\) Engine](#) for all configuration options of the graph server (PGX).

## 14.11.3 Retrieving and Inspecting the Execution Environment

Execution environments are bound to a session. You can retrieve the execution environment for a session by calling `getExecutionEnvironment()` on a `PgxSession`:

- [JShell](#)
- [Java](#)

## JShell

```
opg4j> execEnv.getValues()
==> [analysis-pool.max_num_threads=4, analysis-pool.weight=4, analysis-
pool.priority=MEDIUM, io-pool.num_threads_per_task=4, fast-track-
analysis-pool.max_num_threads=4, fast-track-analysis-pool.weight=1,
fast-track-analysis-pool.priority=HIGH]
```

## Java

```
import oracle.pgx.api.*;
import java.util.List;
import java.util.Map.Entry;

List<Entry<String, Object>> currentValues = execEnv.getValues();
for (Entry<String, Object> value : currentValues) {
 System.out.println(value.getKey() + " = " + value.getValue());
}
```

See [Enterprise Scheduler Configuration Guide](#) for the values of an unmodified execution environment.

To retrieve the sub-environments use the `getIoEnvironment()`, `getAnalysisEnvironment()` and `getFastAnalysisEnvironment()` methods. Each sub-environment has their own `getValues()` method for retrieving the configuration of the sub-environment.

- [JShell](#)
- [Java](#)

## JShell

```
opg4j> var ioEnv = execEnv.getIoEnvironment()
ioEnv ==> IoEnvironment[pool=io-pool]
opg4j> ioEnv.getValues()
$5 ==> {num_threads_per_task=4}

opg4j> var analysisEnv = execEnv.getAnalysisEnvironment()
analysisEnv ==> CpuEnvironment[pool=analysis-pool]
opg4j> analysisEnv.getValues()
$7 ==> {max_num_threads=4, weight=4, priority=MEDIUM}
```

```

opg4j> var fastAnalysisEnv = execEnv.getFastAnalysisEnvironment()
fastAnalysisEnv ==> CpuEnvironment[pool=fast-track-analysis-pool]
opg4j> fastAnalysisEnv.getValues()
$9 ==> {max_num_threads=4, weight=1, priority=HIGH}

```

## Java

```

import oracle.pgx.api.*;
import oracle.pgx.api.executionenvironment.*;
import java.util.Map;

IoEnvironment ioEnv = execEnv.getIoEnvironment();
CpuEnvironment analysisEnv = execEnv.getAnalysisEnvironment();
CpuEnvironment fastAnalysisEnv = execEnv.getFastAnalysisEnvironment();

for (Entry<String, Object> value : ioEnv.getValues().getEntrySet()) {
 System.out.println(value.getKey() + " = " + value.getValue());
}

for (Entry<String, Object> value : analysisEnv.getValues().getEntrySet()) {
 System.out.println(value.getKey() + " = " + value.getValue());
}

for (Entry<String, Object> value :
fastAnalysisEnv.getValues().getEntrySet()) {
 System.out.println(value.getKey() + " = " + value.getValue());
}

```

---

### 14.11.4 Modifying and Submitting Tasks Under an Updated Environment

You can modify an Input/Output (IO) environment in the number of threads by using the `setNumThreadsPerTask()` method of the `IoEnvironment`. The value is updated immediately and all tasks that are submitted after updating it are executed with the updated value.

- [JShell](#)
- [Java](#)

#### JShell

```

opg4j> ioEnv.setNumThreadsPerTask(8)
opg4j> var g = session.readGraphWithProperties(...)
==> PgxGraph[name=graph,N=3,E=6,created=0]

```

## Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.executionenvironment.*;

ioEnv.setNumThreadsPerTask(8);
PgxGraph g = session.readGraphWithProperties(...);
```

---

You can reset an environment to their initial values by calling the `ioEnv.reset()` method. Additionally, you can reset all environments at once by calling `execEnv.reset()` on the `ExecutionEnvironment` class.

You can modify CPU environments in their weight, priority and maximum number of threads using the `setWeight()`, `setPriority()` and `setMaxThreads()` methods:

- 
- [JShell](#)
  - [Java](#)

## JShell

```
opg4j> analysisEnv.setWeight(50)
opg4j> fastAnalysisEnv.setMaxNumThreads(1)
opg4j> var rank = analyst.pagerank(g)
rank ==> VertexProperty[name=pagerank,type=double,graph=my-graph]
```

## Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.executionenvironment.*;

analysisEnv.setWeight(50);
fastAnalysisEnv.setMaxThreads(1);
Analyst analyst = session.createAnalyst();
VertexProperty rank = analyst.pagerank(g);
```

---

### 14.11.5 Using Lambda Syntax

Generally you can perform the following actions in the environment:

1. Set up the execution environment
2. Execute task
3. Reset execution environment

All these actions can be combined and performed in a single step using the `set` method. For each `set` method there is a method using the `with` prefix which takes the updated value and a lambda which should be executed using the updated value.

For example, use `withNumThreadsPerTask()` instead of `setNumThreadsPerTask()` as shown:

- 
- [JShell](#)
  - [Java](#)

## JShell

```
opg4j> var g = ioEnv.withNumThreadsPerTask(8, () ->
session.readGraphWithProperties(...))
==> PgxGraph[name=graph,N=3,E=6,created=0]
```

## Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.executionenvironment.*;

PgxGraph g = ioEnv.withNumThreadsPerTask(8, () ->
session.readGraphWithProperties(...));
```

---

The preceding code execution is equivalent to the following sequence of actions:

```
var oldValue = ioEnv.getNumThreadsPerTask()
ioEnv.setNumThreadsPerTask(currentValue)
var g = session.readGraphWithProperties(...)
ioEnv.setNumThreadsPerTask(oldValue)
```

## 14.12 Admin API

This guide shows how to use the graph server (PGX) Admin API to inspect the server state including sessions, graphs, tasks, memory and thread pools.

- [Get a Server Instance](#)
- [Get Inspection Data](#)
- [Get Active Sessions](#)
- [Get Cached Graphs](#)
- [Get Published Graphs](#)
- [Get Currently Loading Graphs](#)
- [Get Tasks](#)
- [Get Available Memories](#)

## 14.12.1 Get a Server Instance

You can get a PGX Instance as shown in the following code:

- 
- [Java](#)
  - [Python](#)

### Java

```
import oracle.pgx.api.*;
ServerInstance instance = Pgx.getInstance(Pgx.EMBEDDED_URL);
```

### Python

```
instance = pypgx.get_session(base_url = "url")
```

---

## 14.12.2 Get Inspection Data

Inspection data is information about the server state.

You can get the inspection data using the following code:

- 
- [Java](#)
  - [Python](#)

### Java

```
JsonNode serverState = instance.getServerState();
```

### Python

```
server_state = instance.get_server_state()
```

---

This returns a `JsonNode` which contains all administration information, such as, number of graphs loaded, number of sessions, memory usage for graphs and properties and so on.

```
{
 "cached_graphs": [],
```

```

 "published_graphs": [],
 "graphs_currently_loading": [],
 "sessions": [],
 "tasks": [],
 "pools": [],
 "memory": {}
 }

```

### 14.12.3 Get Active Sessions

`serverState.get("sessions")` returns an array of current active sessions. Each entry contains information about a session.

```

{
 "session_id": "530b5f9a-75c4-4838-9cc3-44df44b035c5",
 "source": "testServerState",
 "task_timeout_ms": 0,
 "idle_timeout_ms": 0,
 "alive_ms": 237,
 "total_analysis_time_ms": 115,
 "state": "RELEASED",
 "private_graphs": [
 {
 "name": "anonymous_graph_1",
 "creation_timestamp": 1589317879755,
 "is_transient": true,
 "memory": {
 "topology_bytes": 46,
 "key_mapping_bytes": 30,
 "persistent_property_mem_bytes": 0,
 "transient_property_mem_bytes": 0
 },
 "vertices_num": 1,
 "edges_num": 0,
 "persistent_vertex_properties": [
],
 "persistent_edge_properties": [
],
 "transient_vertex_properties": [
],
 "transient_edge_properties": [
]
 }
],
 "published_graphs": [
 {
 "name": "multigraph",
 "creation_timestamp": 1589317879593,
 "is_transient": false,
 "memory": {

```



```

 "topology_bytes":110,
 "key_mapping_bytes":56,
 "persistent_property_mem_bytes":64,
 "transient_property_mem_bytes":0
 },
 "vertices_num":2,
 "edges_num":6,
 "persistent_vertex_properties":[
 {
 "loaded":true,
 "mem_size_bytes":16,
 "name":"tProp",
 "type":"string"
 }
],
 "persistent_edge_properties":[
 {
 "loaded":true,
 "mem_size_bytes":48,
 "name":"cost",
 "type":"double"
 }
],
 "transient_vertex_properties":[
],
 "transient_edge_properties":[
]
]
}
}

```

The following table explains session information fields:

**Table 14-6 Session Information Options**

Field	Description
sessionID	Session' ID generated by PGX server.
source	Descriptive string identifying the client session.
task_timeout_ms	Timeout to interrupt long-running tasks submitted by sessions (algorithms, I/O tasks) in milliseconds. Set to zero for infinity/no timeout.
idle_timeout_ms	Timeout of idling sessions in milliseconds. Set to zero for infinity/no timeout.
alive_ms	Session's age in milliseconds.
total_analysis_time_ms	Total session's executing time in milliseconds.
state	Current session of the session can be Idle, Submitted, released or terminating.
private_graphs	Session bounded graphs.
published_graphs	Published graphs pointed to from the session.

 **Note:**

The `is_transient` field indicates if the graph is transient. A graph is transient if it is not loaded from an external source.

## 14.12.4 Get Cached Graphs

The server state contains also cached graph information `serverState.get("cached_graphs")` which returns a collection of graphs cached in memory. Each entry contains information about a graph as shown:

```
{
 "name":"sf-1589317879394",
 "creation_timestamp":1589317879394,
 "vertex_properties":[
 {
 "loaded":true,
 "mem_size_bytes":478504,
 "name":"prop1",
 "type":"double"
 }
],
 "edge_properties":[
 {
 "loaded":true,
 "mem_size_bytes":1197720,
 "name":"cost",
 "type":"double"
 },
 {
 "loaded":true,
 "mem_size_bytes":598860,
 "name":"0",
 "type":"integer"
 }
],
 "memory":{
 "topology_bytes":3921814,
 "key_mapping_bytes":1407466,
 "property_mem_bytes":2275084
 },
 "vertices_num":59813,
 "edges_num":149715
}
```

The following table explains graph information fields:

**Table 14-7 Graph Information**

Field	Description
name	Name of the graph.

**Table 14-7 (Cont.) Graph Information**

Field	Description
creation_timestamp	Creation timestamp of the graph.
vertex_properties	List of vertex properties, each entry contains the name, type, memory size used by the property, and a boolean flag to indicate if the property is loaded into memory.
edge_properties	List of edges properties, similar to vertex properties.
memory	Memory size used by the whole graph (topology, key mappings and properties).
vertices_num	Number of vertices.
edges_num	Number of edges.

## 14.12.5 Get Published Graphs

`serverState.get("published_graphs")` returns a list of published graphs.

Each graph entry contains information about the published graph, similar to `cached_graphs`.

## 14.12.6 Get Currently Loading Graphs

`serverState.get("graphs_currently_loading")` returns progress information about graphs which are currently loading.

Each entry, corresponding to one graph, is shown as follows:

```
{
 "name": "anonymous_graph_1",
 "session_id": "530b5f9a-75c4-4838-9cc3-44df44b035c5",
 "start_loading_timestamp": 1605468453030,
 "elapsed_loading_time_ms": 281742,
 "num_vertices_read": 10000000,
 "num_edges_read": 196500000,
 "num_edge_providers_loaded": 1,
 "num_edge_providers_remaining": 9,
 "num_vertex_providers_loaded": 1,
 "num_vertex_providers_remaining": 0,
 "loading_phase": "reading edges",
 "loading_phase_start_timestamp": 1605468453085,
 "loading_phase_elapsed_time_ms": 281687,
 "loading_phase_state": "current vertex provider index: 1, number
of vertices read for prorvider: 0, current edge provider index: 1,
number of edges read for prorvider: 76,500,000"
}
```

The `name` field contains a temporary name of the graph. It may not be equal to the name that is assigned to graph after loading.

Fields indicating the number of read vertices and edges are updated in regular intervals of 10,000 entities.

The field `loading_phase` indicates the current phase during graph loading. Valid values are "reading edges" or "building graph indices". For some loading phases, the field `loading_phase_state` contains a string with additional information on the phase. However, not all loading phases provide this additional information.



#### Note:

`graphs_currently_loading` is supported for data formats CSV, ADJ\_LIST, EDGE\_LIST, TWO\_TABLES and PG (FLAT\_FILE) for homogeneous graphs and for formats CSV and RDBMS for partitioned graphs.

## 14.12.7 Get Tasks

`serverState.get("tasks")` returns the last 100 queued tasks.

Each task has a `type`, the pool to be executed on (the task might be already executed) and other status fields (`{Queued|Started|Done}` time), and a `sessionid` if the task belongs to a session.

## 14.12.8 Get Available Memories

This section contains a map of available memories, the key is the hostname and the value is a list of current available memories (managed and unmanaged). Each entry contains how much memory is free, used and the maximum available memory.

## 14.13 PgxFrames Tabular Data-Structure

`PgxFrame` is a data-structure to load, store and manipulate tabular data. It contains rows and columns. A `PgxFrame` can contain multiple columns where each column consist of elements of the same data type, and has a name. The list of the columns with their names and data types defines the schema of the frame. (The number of rows in the `PgxFrame` is not part of the schema of the frame.)

`PgxFrame` provides some operations that also output `PgxFrames` (described later in the tutorial). Those operations can be performed in-place (meaning that the frame is mutated during the operation) in order to save memory. In place operations should be used whenever possible. However, we provide out-place variants, i.e., a new frame is created during the operation.

The following table lists all the in-place operations along with the respective out-place operations:

**Table 14-8 Mapping between In-Place and Out-Place Operations**

In-place operations	Out-place operations
<code>headInPlace</code>	<code>head</code>
<code>tailInPlace</code>	<code>tail</code>
<code>flattenAllInPlace</code>	<code>flattenAll</code>
<code>renameColumnInPlace</code>	<code>renameColumn</code>
<code>renameColumnsInPlace</code>	<code>renameColumns</code>

**Table 14-8 (Cont.) Mapping between In-Place and Out-Place Operations**

In-place operations	Out-place operations
selectInPlace	select

- [Loading a PgxFFrame from a Database](#)
- [Loading a PgxFFrame from Client-Side Data](#)
- [Printing the Content of a PgxFFrame](#)
- [Destroying a PgxFFrame](#)
- [Storing a PgxFFrame to a Database](#)
- [Loading and Storing Vector Properties](#)
- [Flattening Vector Properties](#)
- [Union of PGX Frames](#)
- [Joining PGX Frames](#)
- [PgxFFrame Helpers](#)
- [PgxFFrame-PgqlResultSet Conversions](#)
- [Creating a Graph from Multiple PgxFFrame Objects](#)

### 14.13.1 Loading a PgxFFrame from a Database

`PgxFFrame(S)` can also be loaded from relational tables in an Oracle database. Each column of the relational table will correspond to a column in the loaded frame. When loading `PgxFrames` from the database, the default behavior is to detect what columns the table has, and to load them all. If not specified explicitly, the connection details of the current user and session are used and the columns are detected automatically. The following describes the steps to load a `PgxFFrame` from a database table:

1. Create a **Session** and an **Analyst**:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
cd /opt/oracle/graph/
./bin/opg4j
// starting the shell will create an implicit session and analyst
opg4j> import static
oracle.pgx.api.frames.functions.ColumnRenaming.renaming
opg4j> import static
oracle.pgx.api.frames.schema.ColumnDescriptor.columnDescriptor
```

```
opg4j> import oracle.pgx.api.frames.schema.*
opg4j> import oracle.pgx.api.frames.schema.datatypes.*
```

## Java

```
import oracle.pgx.api.*;
import oracle.pgx.api.frames.*;
import oracle.pgx.api.frames.functions.*;
import oracle.pgx.api.frames.schema.*;
import oracle.pgx.api.frames.schema.datatypes.*;
import static oracle.pgx.api.frames.functions.ColumnRenaming.renaming;
import static
oracle.pgx.api.frames.schema.ColumnDescriptor.columnDescriptor;

PgxSession session = Pgx.createSession("my-session");
Analyst analyst = session.createAnalyst();
```

## Python

```
session = pypgx.get_session(session_name="my-session")
analyst = session.create_analyst()
```

---

## 2. Load a PgxFrames.

---

- [JShell](#)
- [Java](#)

## JShell

```
opg4j> var exampleFrame = session.readFrame().
 db().
 name("framename"). // name of the frame
 tablename("tablename"). // name of the table from where the data must
be loaded
 connections(16). // indicates that 16 connections can be used
to load in parallel
 load()
```

## Java

```
PgxFrames exampleFrame = session.readFrame()
 .db()
 .name("framename") // name of the frame
 .tablename("tablename") // name of the table from where the data must
be loaded
 .connections(16) // indicates that 16 connections can be used
```

```
to load in parallel
 .load();
```

- 
3. If only a subset of the columns must be loaded, then you must specify the columns with `FrameReader.columns()`.
- 

- [JShell](#)
- [Java](#)

## JShell

```
// You must specify jdbc connection, keystore and the columns to
load
opg4j> session.registerKeystore("keystore", pathToKeystore,
keystorePassword)
opg4j> var exampleFrame = session.readFrame().
 db().
 name("framename").
 tablename("tablename"). // name of the table from where
the data must be loaded
 jdbcUrl("jdbcUrl").
 username("user").
 keystoreAlias("keytores").
 owner("owner"). // necessary if the table is owned
by another user
 connections(16). // indicates that 16 connections
can be used to load in parallel
 columns(exampleFrameSchema). // columns to load
 load()
```

## Java

```
import oracle.pgx.api.frames.schema.datatypes.DataTypes;
import oracle.pgx.api.frames.schema.ColumnDescriptor;
// You must specify jdbc connection, keystore and the columns to
load
session.registerKeystore("keystore", pathToKeystore,
keystorePassword)
PgxFrame exampleFrame = session.readFrame()
 .db()
 .name("framename")
 .tablename("tablename") // name of the table from where
the data must be loaded
 .jdbcUrl("jdbcUrl")
 .username("user")
 .keystoreAlias("keytores")
 .owner("owner") // necessary if the table is owned
by another user
```

```
.connections(16) // indicates that 16 connections can be
used to load in parallel
.columns(exampleFrameSchema) // columns to load
.load();
```

---

You can also create a graph from the `PgxFrames`(s). See [Creating a Graph from Multiple PgxFrames Objects](#) for more information.

## 14.13.2 Loading a PgxFrames from Client-Side Data

You can also load `PgxFrames`(s) directly from client-side data.

The following describes the steps to load a `PgxFrames` from client-side data:

1. Create a **Session** and an **Analyst**:

See step-1 in [Loading a PgxFrames from a Database](#) for the code examples.

2. Define a frame schema to load a `PgxFrames` from client side data. For example, the following shows a frame schema defined with various data types:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
opg4j> var exampleFrameSchema = List.of(
 columnDescriptor("name", DataTypes.STRING_TYPE),
 columnDescriptor("age", DataTypes.INTEGER_TYPE),
 columnDescriptor("salary", DataTypes.DOUBLE_TYPE),
 columnDescriptor("married", DataTypes.BOOLEAN_TYPE),
 columnDescriptor("tax_rate", DataTypes.FLOAT_TYPE),
 columnDescriptor("random", DataTypes.LONG_TYPE),
 columnDescriptor("date_of_birth", DataTypes.LOCAL_DATE_TYPE)
)
```

### Java

```
List<ColumnDescriptor> exampleFrameSchema = Arrays.asList(
 columnDescriptor("name", DataTypes.STRING_TYPE),
 columnDescriptor("age", DataTypes.INTEGER_TYPE),
 columnDescriptor("salary", DataTypes.DOUBLE_TYPE),
 columnDescriptor("married", DataTypes.BOOLEAN_TYPE),
 columnDescriptor("tax_rate", DataTypes.FLOAT_TYPE),
 columnDescriptor("random", DataTypes.LONG_TYPE),
 columnDescriptor("date_of_birth", DataTypes.LOCAL_DATE_TYPE)
);
```



## Python

```
example_frame_schema = [
 ("name", "STRING_TYPE"),
 ("age", "INTEGER_TYPE"),
 ("salary", "DOUBLE_TYPE"),
 ("married", "BOOLEAN_TYPE"),
 ("tax_rate", "FLOAT_TYPE"),
 ("random", "LONG_TYPE"),
 ("date_of_birth", "LOCAL_DATE_TYPE")
]
```

---

### 3. Define data as per the schema.

---

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> Map<String, Iterable<?>> exampleFrameData = Map.of(
 "name", Arrays.asList("Alice", "Bob", "Charlie"),
 "age", Arrays.asList(25, 27, 29),
 "salary", Arrays.asList(10000.0, 15000.0, 20000.0),
 "married", Arrays.asList(false, false, true),
 "tax_rate", Arrays.asList(0.21, 0.26, 0.32),
 "random", Arrays.asList(2394293898324L, 45640604960495L,
12312323409087654L),
 "date_of_birth", Arrays.asList(
 LocalDate.of(1990, 9, 15),
 LocalDate.of(1991, 11, 4),
 LocalDate.of(1993, 10, 4)
)
)
```

## Java

```
Map<String, Iterable<?>> exampleFrameData = new HashMap<>();
exampleFrameData.put("name", Arrays.asList("Alice", "Bob",
"Charlie"));
exampleFrameData.put("age", Arrays.asList(25, 27, 29));
exampleFrameData.put("salary", Arrays.asList(10000.0, 15000.0,
20000.0));
exampleFrameData.put("married", Arrays.asList(false, false, true));
exampleFrameData.put("tax_rate", Arrays.asList(0.21, 0.26, 0.32));
exampleFrameData.put("random", Arrays.asList(2394293898324L,
45640604960495L, 12312323409087654L));
exampleFrameData.put("date_of_birth",
```

```
 Arrays.asList(LocalDate.of(1990, 9, 15),
 LocalDate.of(1991, 11, 4),
 LocalDate.of(1993, 10, 4)
)
);
```

## Python

```
from datetime import date

example_frame_data = {
 "name": ["Alice", "Bob", "Charlie"],
 "age": [25, 27, 29],
 "salary": [10000.0, 15000.0, 20000.0],
 "married": [False, False, True],
 "tax_rate": [0.21, 0.26, 0.32],
 "random": [2394293898324, 45640604960495, 12312323409087654],
 "date_of_birth": [date(1990, 9, 15),
 date(1991, 11, 4),
 date(1993, 10, 4)]
}
```

---

#### 4. Load the frame as shown:

---

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var exampleFrame = session.createFrame(exampleFrameSchema,
exampleFrameData, "example frame")
```

## Java

```
PgxFrame exampleFrame = session.createFrame(exampleFrameSchema,
exampleFrameData, "example frame");
```

## Python

```
example_frame=session.create_frame(example_frame_schema,example_frame_data
, 'example frame')
```

---

#### 5. You can also load the frame incrementally as you receive more data:

---

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var exampleFrameBuilder =
session.createFrameBuilder(exampleFrameSchema);
opg4j> exampleFrameBuilder.addRow(exampleFrameData)
opg4j> Map<String, Iterable<?>> exampleFrameDataPart2 = Map.of(
 "name", Arrays.asList("Dave"),
 "age", Arrays.asList(26),
 "salary", Arrays.asList(18000.0),
 "married", Arrays.asList(true),
 "tax_rate", Arrays.asList(0.30),
 "random", Arrays.asList(456783423423L),
 "date_of_birth", Arrays.asList(LocalDate.of(1989, 9, 15))
)
opg4j> exampleFrameBuilder.addRow(exampleFrameDataPart2)
opg4j> var exampleFrame = exampleFrameBuilder.build("example frame")
```

## Java

```
PgxFrameBuilder exampleFrameBuilder =
session.createFrameBuilder(exampleFrameSchema);
exampleFrameBuilder.addRow(exampleFrameData);
Map<String, Iterable<?>> exampleFrameDataPart2 = new HashMap<>();
exampleFrameDataPart2.put("name", Arrays.asList("Dave"));
exampleFrameDataPart2.put("age", Arrays.asList(26));
exampleFrameDataPart2.put("salary", Arrays.asList(18000.0));
exampleFrameDataPart2.put("married", Arrays.asList(true));
exampleFrameDataPart2.put("tax_rate", Arrays.asList(0.30));
exampleFrameDataPart2.put("random", Arrays.asList(456783423423L));
exampleFrameDataPart2.put("date_of_birth",
 Arrays.asList(LocalDate.of(1989, 9, 15))
);
exampleFrameBuilder.addRow(exampleFrameDataPart2);
PgxFrame exampleFrame = exampleFrameBuilder.build("example frame");
```

## Python

```
example_frame_builder =
session.create_frame_builder(example_frame_schema)
example_frame_builder.add_rows(example_frame_data)
example_frame_data_part_2 = {
 "name": ["Dave"],
 "age": [26],
 "salary": [18000.0],
 "married": [True],
```

```

 "tax_rate": [0.30],
 "random": [456783423423],
 "date_of_birth": [date(1989, 9, 15)]
}
example_frame_builder.add_rows(example_frame_data_part_2)
example_frame = example_frame_builder.build("example frame")

```

- 
6. Finally, you can also load a frame from a Pandas dataframe in Python as shown:

```

import pandas as pd
example_pandas_dataframe = pd.DataFrame(data=example_frame_data)
example_frame = session.pandas_to_pgx_frame(example_pandas_dataframe,
"example frame")

```

You can also create a graph from the `PgxFrame(s)`. See [Creating a Graph from Multiple PgxFrame Objects](#) for more information.

### 14.13.3 Printing the Content of a PgxFrame

You can observe the contents of a frame using the `print` functionality as shown:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> exampleFrame.print()
```

#### Java

```
exampleFrame.print();
```

#### Python

```
example_frame.print()
```

---

The output appears as follows:

```

+-----+
----+
| name | age | salary | married | tax_rate | random |
date_of_birth |

```

```
+-----+
-----+
| John | 27 | 4133300.0 | true | 11.0 | 123456782 |
1985-10-18 |
| Albert | 23 | 5813000.5 | false | 12.0 | 124343142 |
2000-01-14 |
| Heather | 28 | 1.0130302E7 | true | 10.5 | 827520917 |
1985-10-18 |
| Emily | 24 | 9380080.5 | false | 13.0 | 128973221 |
1910-07-30 |
| "D'Juan" | 27 | 1582093.0 | true | 11.0 | 92384 |
1955-12-01 |
+-----+
-----+
```

### 14.13.4 Destroying a PgxFFrame

`PgxFrames` consumes a lot of memory on the graph server (PGX) if they have a lot of rows or columns. Hence it is necessary to close them with the `close()` operation. After this operation, the content of the `PgxFFrame` is not available anymore.

You can close a frame as shown:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> exampleFrame.close()
```

#### Java

```
exampleFrame.close();
```

#### Python

```
exampleFrame.close()
```

---

### 14.13.5 Storing a PgxFFrame to a Database

When storing a `PgxFFrame` into the database, the frame is stored as a table, where the columns correspond to the columns of the `PgxFFrame` and the rows correspond to the rows of the `PgxFFrame`.

**Note:**

Column order preservation may or may not happen when storing a `PgxFFrame` in the database.

**Overwrite Mode**

When storing as a table in the database, you can overwrite an already existing table (with correct structure).

In overwrite mode, the previous table is truncated (emptied), and the data is then inserted. By default, it is set to `false` so that if a table already exists, it will throw an error to the user unless `overwrite` is set to `true`.

- [JShell](#)
- [Java](#)

**JShell**

```
// store as table in the database using jdbc + username + password
opg4j> exampleFrame.write().
 db(). // select the "format" to be relational db
 name("framename"). // name of the frame
 tablename("tablename"). // name of the table in which the data must be
stored
 overwrite(true). // indicates that if there is a table with the
same name, it will be overwritten (truncated)
 connections(16). // indicates that 16 connections can be used
to store in parallel
 jdbcUrl("jdbcUrl").
 username("user").
 password("password").
 store()
```

**Java**

```
exampleFrame.write()
 .db() // select the "format" to be relational db
 .name("framename") // name of the frame
 .tablename("tablename") // name of the table in which the data must be
stored
 .overwrite(true) // indicates that if there is a table with the
same name, it will be overwritten (truncated)
 .connections(16) // indicates that 16 connections can be used
to store in parallel
 .jdbcUrl("jdbcUrl")
 .username("user")
 .password("password")
 .store();
```

---

## 14.13.6 Loading and Storing Vector Properties

You can load or store vector properties which are fundamental for PgxML functionality in the graph server (PGX) using PgxFrames.

In order to load a `PgxFrame` with vector properties, follow the steps as shown:

1. Create the `PgxFrame` schema, defining the columns as shown:

- 
- [JShell](#)
  - [Java](#)

### JShell

```
opg4j> var vecFrameSchema = List.of(
 columnDescriptor("intProp", DataTypes.INTEGER_TYPE),
 columnDescriptor("intProp2", DataTypes.INTEGER_TYPE),
 columnDescriptor("vectProp",
DataTypes.vector(DataTypes.FLOAT_TYPE, 3)),
 columnDescriptor("stringProp", DataTypes.STRING_TYPE),
 columnDescriptor("vectProp2",
DataTypes.vector(DataTypes.FLOAT_TYPE, 2))
).toArray(new ColumnDescriptor[0])
```

### Java

```
ColumnDescriptor[] vecFrameSchema = {
 columnDescriptor("intProp", DataTypes.INTEGER_TYPE),
 columnDescriptor("intProp2", DataTypes.INTEGER_TYPE),
 columnDescriptor("vectProp",
DataTypes.vector(DataTypes.FLOAT_TYPE, 3)),
 columnDescriptor("stringProp", DataTypes.STRING_TYPE),
 columnDescriptor("vectProp2",
DataTypes.vector(DataTypes.FLOAT_TYPE, 2))
};
```

- 
2. Load the `PgxFrame` with the given schema from the specified path:

- 
- [JShell](#)
  - [Java](#)

## JShell

```

opg4j> var vecFrame = session.readFrame().
 db().
 name("vector PgxFFrame").
 tablename("tablename"). // name of the table from where the data
must be loaded
 jdbcUrl("jdbcUrl").
 username("user").
 owner("owner"). // necessary if the table is owned by
another user
 connections(16). // indicates that 16 connections can be
used to load in parallel
 columns(vecFrameSchema). // columns to load
 load()

```

## Java

```

PgxFFrame vecFrame = session.readFrame()
 .db()
 .name("vector PgxFFrame")
 .tablename("tablename") // name of the table from where the data
must be loaded
 .jdbcUrl("jdbcUrl")
 .username("user")
 .owner("owner") // necessary if the table is owned by
another user
 .connections(16) // indicates that 16 connections can be
used to load in parallel
 .columns(vecFrameSchema) // columns to load
 .load();

```

The final result in the `PgxFFrame` may appear as follows:

```

+-----+
| intProp | intProp2 | vectProp | stringProp | vectProp2 |
+-----+
| 0 | 2 | 0.1;0.2;0.3 | testProp0 | 0.1;0.2 |
| 1 | 1 | 0.1;0.2;0.3 | testProp10 | 0.1;0.2 |
| 1 | 2 | 0.1;0.2;0.3 | testProp20 | 0.1;0.2 |
| 2 | 3 | 0.1;0.2;0.3 | testProp30 | 0.1;0.2 |
| 3 | 1 | 0.1;0.2;0.3 | testProp40 | 0.1;0.2 |
+-----+

```

### 14.13.7 Flattening Vector Properties

You can split the vector properties into multiple columns using the `flattenAll()` operation.

For example, you can flatten the vector properties for the example explained in [Loading and Storing Vector Properties](#) as shown:



- [JShell](#)
- [Java](#)

## JShell

```
opg4j> vecFrame.flattenAll()
```

## Java

```
vecFrame.flattenAll();
```

The resulting flattened `PgxFrame` may appear as shown:

```
+-----+
+-----+
| intProp | intProp2 | vectProp_0 | vectProp_1 | vectProp_2 |
stringProp | vectProp2_0 | vectProp2_1 |
+-----+
+-----+
| 0 | 2 | 0.1 | 0.2 | 0.3 |
testProp0 | 0.1 | 0.2 | | |
| 1 | 1 | 0.1 | 0.2 | 0.3 |
testProp10 | 0.1 | 0.2 | | |
| 1 | 2 | 0.1 | 0.2 | 0.3 |
testProp20 | 0.1 | 0.2 | | |
| 2 | 3 | 0.1 | 0.2 | 0.3 |
testProp30 | 0.1 | 0.2 | | |
| 3 | 1 | 0.1 | 0.2 | 0.3 |
testProp40 | 0.1 | 0.2 | | |
+-----+
+-----+
```

### 14.13.8 Union of PGX Frames

You can join two `PgxFrames` that have compatible columns (i.e. same type and order).

- [JShell](#)
- [Java](#)

## JShell

```
opg4j> <first-frame>.union(<secondframe>).print()
```

## Java

```
<first-frame>.union(<secondframe>).print();
```

The rows of the resulting `PgxFrame` are the union of the rows from the two original frames.

**Note:**

The union operation will not remove duplicate rows that resulted from the union operation.

### 14.13.9 Joining PGX Frames

You can join two frames whose rows are correlated through one of the columns using the `join` functionality. This allows us to combine frames by checking for equality between rows for a specific column.

The following example shows joining two `PgxFrames` `exampleFrame` and `moreInfoFrame` on the `name` column by calling the `join` method.

- [JShell](#)
- [Java](#)

#### JShell

```
opg4j> exampleFrame.join(moreInfoFrame, "name", "leftFrame",
"rightFrame").print()
```

#### Java

```
exampleFrame.join(moreInfoFrame, "name", "leftFrame", "rightFrame").print();
```

The result may appear as shown:

```
+-----+
+-----+
+-----+
| leftFrame_name | leftFrame_age | leftFrame_salary | leftFrame_married |
| leftFrame_tax_rate | leftFrame_random | leftFrame_date_of_birth |
| rightFrame_name | rightFrame_title | rightFrame_reports |
+-----+
+-----+
```

```

-----+
| John | 27 | 4133300.0 | |
true | 11.0 | 123456782 | |
1985-10-18 | John | Software Engineering
Manager | 5
| Albert | 23 | 5813000.5 | |
false | 12.0 | 124343142 | |
2000-01-14 | Albert | Sales
Manager | 10
| Emily | 24 | 9380080.5 | |
false | 13.0 | 128973221 | |
1910-07-30 | Emily | Operations
Manager | 20
-----+
-----+
-----+

```

The joined frame contains the columns of the two frames involved in the operation for the rows with the same `name`.



#### Note:

The column prefixes specified in the `join()` call, `leftFrame` and `rightFrame`.

## 14.13.10 PgxFrame Helpers

`PgxFrame` supports the following operations:

- `head`
- `tail`
- `select`
- `renameColumns`

### Head Operation

The `head` operation can be used to only keep the first rows of a `PgxFrame`. (The result is deterministic only for ordered `PgxFrame`.)

- [JShell](#)
- [Java](#)

### JShell

```
opg4j> vecFrame.head(2).print()
```

## Java

```
vecFrame.head(2).print();
```

The output appears as follows:

```

+-----+
| intProp | intProp2 | vectProp | stringProp | vectProp2 |
+-----+
| 0 | 2 | 0.1;0.2;0.3 | testProp0 | 0.1;0.2 |
| 1 | 1 | 0.1;0.2;0.3 | testProp10 | 0.1;0.2 |
+-----+

```

### Tail Operation

The `tail` operation can be used to only keep the last rows of a `PgxFrame`. (The result is deterministic only for ordered `PgxFrame`).

- [JShell](#)
- [Java](#)

## JShell

```
opg4j> vecFrame.tail(2).print()
```

## Java

```
vecFrame.tail(2).print();
```

The output appears as follows:

```

+-----+
| intProp | intProp2 | vectProp | stringProp | vectProp2 |
+-----+
| 2 | 3 | 0.1;0.2;0.3 | testProp30 | 0.1;0.2 |
| 3 | 1 | 0.1;0.2;0.3 | testProp40 | 0.1;0.2 |
+-----+

```

### Select Operation

The `select` operation can be used to keep only a specified list of columns of an input `PgxFrame`.

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var vecFrameSelected = vecFrame.select("vectProp2", "vectProp",
"stringProp")
```

## Java

```
PgxFrame vecFrameSelected =
vecFrame.select("vectProp2", "vectProp", "stringProp");
```

## Python

```
vec_frame_selected =
vec_frame.select("vectProp2", "vectProp", "stringProp")
```

---

The result may appear as follows:

```
+-----+
| vectProp2 | vectProp | stringProp |
+-----+
| 0.1;0.2 | 0.1;0.2;0.3 | testProp0 |
| 0.1;0.2 | 0.1;0.2;0.3 | testProp10 |
| 0.1;0.2 | 0.1;0.2;0.3 | testProp20 |
| 0.1;0.2 | 0.1;0.2;0.3 | testProp30 |
| 0.1;0.2 | 0.1;0.2;0.3 | testProp40 |
+-----+
```

## Rename PgxFrame Columns

You can rename the columns in a `PgxFrame` to customized names as follows:

- 
- [JShell](#)
  - [Java](#)

## JShell

```
opg4j> var vecFrameRenamed = vecFrame.renameColumns(
 renaming("vectProp2", "vectProp2_renamed"),
 renaming("vectProp", "vectProp_renamed"),
```

```
 renaming("stringProp", "stringProp_renamed")
)
```

## Java

```
vecFrameRenamed = vecFrame.renameColumns(renaming("vectProp2",
"vectProp2_renamed"),
 renaming("vectProp",
"vectProp_renamed"),
 renaming("stringProp",
"stringProp_renamed"));
```

The renamed `PgxFrame` appears as follows:

```
+-----+
----+
| intProp | intProp2 | vectProp_renamed | stringProp_renamed |
vectProp2_renamed |
+-----+
----+
| 0 | 2 | 0.1;0.2;0.3 | testProp0 |
0.1;0.2 | | | |
| 1 | 1 | 0.1;0.2;0.3 | testProp10 |
0.1;0.2 | | | |
| 1 | 2 | 0.1;0.2;0.3 | testProp20 |
0.1;0.2 | | | |
| 2 | 3 | 0.1;0.2;0.3 | testProp30 |
0.1;0.2 | | | |
| 3 | 1 | 0.1;0.2;0.3 | testProp40 |
0.1;0.2 | | | |
+-----+
----+
```

### 14.13.11 PgxFrame-PgsqlResultSet Conversions

You can perform conversions between `PgxFrames` and `PgsqlResultSets`.

#### PgxFrame to PgsqlResultSet

You can convert a `PgxFrame` to `PgsqlResultSet` as follows:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var resultSet = exampleFrame.toPgqlResultSet()
```

## Java

```
PgqlResultSet resultSet = exampleFrame.toPgqlResultSet();
```

## Python

```
result_set = example_frame.to_pgql_result_set()
```

You can view the content of the result set through the usual `PgqlResultSet` APIs. The output appears as follows:

```
+-----+
+-----+
| name | age | salary | married | tax_rate | random |
| date_of_birth |
+-----+
+-----+
| John | 27 | 4133300.0 | true | 11.0 | 123456782 |
| 1985-10-18 |
| Albert | 23 | 5813000.5 | false | 12.0 | 124343142 |
| 2000-01-14 |
| Heather | 28 | 1.0130302E7 | true | 10.5 | 827520917 |
| 1985-10-18 |
| Emily | 24 | 9380080.5 | false | 13.0 | 128973221 |
| 1910-07-30 |
| "D'Juan" | 27 | 1582093.0 | true | 11.0 | 92384 |
| 1955-12-01 |
+-----+
+-----+
```

### PgqlResultSet to PgxFrame

You can convert a `PgqlResultSet` to a `PgxFrame` as follows:

- [JShell](#)
- [Java](#)

## JShell

```
opg4j> var query = ...
opg4j> var graph = ...
```

```
opg4j> var resultSet = graph.queryPgql(query)
opg4j> resultSet.toFrame()
```

## Java

```
String query = ...;
PgxGraph graph = ...;
PgqlResultSet resultSet = graph.queryPgql(query);
resultSet.toFrame();
```

### 14.13.12 Creating a Graph from Multiple PgxFrames Objects

You can create a `PgxGraph` with vertex `PgxFrame(s)` and edge `PgxFrame(s)`.

Consider the following `PgxFrame` objects:

```
people
+-----+
| id | name |
+-----+
| 1 | Alice |
| 2 | Bob |
| 3 | Charlie |
+-----+

houses
+-----+
| identification | location |
+-----+
| 1 | Road 1 |
| 2 | Street 5 |
| 3 | Avenue 4 |
+-----+

knows
+-----+
| src | dst |
+-----+
| 1 | 1 |
| 2 | 3 |
| 3 | 2 |
+-----+

lives
+-----+
| source | destination |
+-----+
| 1 | 2 |
| 2 | 1 |
```



```
| 3 | 3 |
+-----+
```

You can now create a `PgxGraph` as shown in the following examples:

- [JShell](#)
- [Java](#)
- [Python](#)

## JShell

```
opg4j> var graphFromFramesCreator =
session.createGraphFromFrames("example graph")
opg4j> graphFromFramesCreator.vertexProvider("people", people)
opg4j> graphFromFramesCreator.vertexProvider("houses",
houses).vertexKeyColumn("identification")
opg4j> graphFromFramesCreator.edgeProvider("knows", "people",
"people", knows)
opg4j> var edge_provider =
graphFromFramesCreator.edgeProvider("lives", "people", "houses", lives)
opg4j> edge_provider.sourceVertexKeyColumn("source")
opg4j> edge_provider.destinationVertexKeyColumn("destination")
opg4j> graphFromFramesCreator.partitioned(true)
opg4j> var graph = graphFromFramesCreator.create()
```

## Java

```
PgxGraphFromFramesCreator graphFromFramesCreator =
session.createGraphFromFrames("example graph");
graphFromFramesCreator.vertexProvider("people", people);
graphFromFramesCreator.vertexProvider("houses",
houses).vertexKeyColumn("identification");
graphFromFramesCreator.edgeProvider("knows", "people", "people",
knows);
PgxEdgeProviderFromFramesCreator edgeProvider =
graphFromFramesCreator.edgeProvider("lives", "people", "houses",
lives);
edgeProvider.sourceVertexKeyColumn("source");
edgeProvider.destinationVertexKeyColumn("destination");
graphFromFramesCreator.partitioned(true);
PgxGraph graph = graphFromFramesCreator.create();
```

## Python

```
vertex_providers_from_frames = [
 session.vertex_provider_from_frame("person",
 people),
 session.vertex_provider_from_frame("house",
```

```
 frame = houses,
 vertex_key_column = "identification")
]
edge_providers_from_frames = [
 session.edge_provider_from_frame("person_knows_person",
 source_provider = "person",
 destination_provider = "person",
 frame = knows),
 session.edge_provider_from_frame("person_lives_at_house",
 source_provider = "person",
 destination_provider = "house",
 frame = lives,
 source_vertex_column="source",
 destination_vertex_column="destination")
]
graph = session.graph_from_frames("example graph",
vertex_providers_from_frames, edge_providers_from_frames, partitioned=True)
```

---

# 15

## Working with Files Using the Graph Server (PGX)

This chapter describes in detail about working with different file formats to perform various actions like loading, storing or exporting a graph using the Graph Server (PGX).

- [Loading Graph Data from Files](#)
- [Loading Graph Data in Parallel from Multiple Files](#)
- [Exporting Graphs Into a File](#)
- [Exporting a Graph into Multiple Files](#)

### 15.1 Loading Graph Data from Files

You can load graph data from files by either of the two ways:

- using the header format specified in the files
- by directly calling the graph builder API

#### Creating a graph using file header format

The graph server (PGX) uses the header of the files to determine the name and types of the properties to load. It also infers the column to be used as vertex ID, the columns that indicate the source and destination vertex ID for edges, and the column to be loaded as vertex or edge label.

#### Creating a graph using graph builder API

You can also use `PgxSession.readGraphFiles()` to load the graph. This method takes the following three arguments:

- path to the vertex file
- path to the edge file
- name of the graph to be created

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

#### JShell

```
opg4j> var loadedGraph = session.readGraphFiles("<path/vertices.csv>", "<path/edges.csv>", "<graph_name>")
```

## Java

```
import oracle.pgx.api.PgxSession;
import oracle.pgx.api.PgxGraph;

PgxSession session = Pgx.createSession("NewSession");
PgxGraph loadedGraph = session.readGraphFiles("<path/vertices.csv>",
"<path/edges.csv>", "<graph_name>");
```

## Python

```
session = pypgx.get_session(session_name="<session_name>")
loaded_graph = session.read_graph_files("<path/vertices.csv>", "<path/edges.csv>", "<graph_name>")
```

The graph server (PGX) supports loading graph data from files for the following data formats

- Plain Text Formats
- XML File Formats
- Binary File Formats
- [Graph Configuration for Loading from File](#)
- [Specifying the File Path](#)
- [Supported File Access Protocols](#)
- [Plain Text Formats](#)
- [XML File Formats](#)
- [Binary File Formats](#)

### 15.1.1 Graph Configuration for Loading from File

The following table presents the graph configuration options to load graph data from all supported file formats to the graph server (PGX).

**Table 15-1 Loading from File - Graph Configuration Options**

Field	Type	Description	Default
array_compaction_threshold	number	<i>[only relevant if the graph is optimized for updates]</i> Threshold used to determine when to compact the delta-logs into a new array. If lower than the engine min_array_compaction_threshold value, min_array_compaction_threshold will be used instead.	0.2
attributes	object	Additional attributes needed to read and write the graph data.	null

**Table 15-1 (Cont.) Loading from File - Graph Configuration Options**

Field	Type	Description	Default
detect_gzip	boolean	Enable or disable automatic gzip compression detection when loading graphs.	true
edge_id_strategy	enum[no_ids, keys_as_ids, unstable_generated_ids]	Indicates what ID strategy should be used for the edges of this graph. If not specified (or set to null), the strategy will be determined during loading or using a default value.	null
edge_id_type	enum[long]	Type of the edge ID. For homogeneous graphs, if not specified (or set to null), it will default to long.	null
edge_props	array of object	Specification of edge properties associated with graph.	[]
edge_uris	array of string	List of unified resource identifiers.	[]
error_handling	object	Error handling configuration.	null
external_stores	array of object	Specification of the external stores where external string properties reside.	[]
format	enum[pgb, edge_list, adj_list, graphml, pg, rdf, two_tables]	Graph format to be used.	null
header	boolean	First line of file is meant for headers. For example, 'EdgeId, SourceId, DestId, EdgeProp1, EdgeProp2'	false
keystore_alias	string	Alias to the keystore to use when connecting to the database.	null
loading	object	Loading-specific configuration.	null
local_date_formats	array of string	Array of local_date formats to use when loading and storing local_date properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string.	[]
optimized_for	enum[read, updates]	Indicates if the graph must use data-structures optimized for read-intensive scenarios or for fast updates.	read
partition_while_loading	enum[by_label, no]	Indicates if the graph must be partitioned while loading.	null
password	string	Password to use when connecting to database.	null
point2d	string	Longitude and latitude as floating point values separated by a space.	0.0 0.0

**Table 15-1 (Cont.) Loading from File - Graph Configuration Options**

Field	Type	Description	Default
separator	string	A series of single-character separators for tokenizing. The characters ", {, } and \n cannot be used as separators. Default value is ", " for CSV files, and "\t " for other formats. The first character will be used as a separator when storing.	null
storing	object	Storing-specific configuration.	null
time_format	array of string	The time format to use when loading and storing time properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string.	[]
time_with_timezone_one_format	array of string	The time with timezone format to use when loading and storing time with timezone properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string.	[]
timestamp_format	array of string	The timestamp format to use when loading and storing timestamp properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string.	[]
timestamp_with_timezone_format	array of string	The timestamp with timezone format to use when loading and storing timestamp with timezone properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string.	[]
vector_component_delimiter	character	Delimiter for the different components of vector properties.	;
vertex_id_strategy	enum[no_ids, keys_as_ids, unstable_generated_ids]	Indicates what ID strategy should be used for the vertices of this graph. If not specified (or set to null), the strategy will be automatically detected.	null
vertex_id_type	enum[int, integer, long, string]	Type of the vertex ID. For homogeneous graphs, if not specified (or set to null), it will default to a specific value (depending on the origin of the data).	null
vertex_props	array of object	Specification of vertex properties associated with graph.	[]
vertex_uris	array of string	List of unified resource identifiers.	[]

In the CSV format, the columns used to specify the `vertex ID column`, `vertex labels column`, `edge ID column`, `edge source ID column`, `edge destination ID column` and the `edge label column` can be configured with the CSV specific fields as shown in the following table:

Table 15-2 CSV Specific Options

Field	Type	Description	Default
array_compaction_threshold	number	<i>[only relevant if the graph is optimized for updates]</i> Threshold used to determined when to compact the delta-logs into a new array. If lower than the engine <code>min_array_compaction_threshold</code> value, <code>min_array_compaction_threshold</code> will be used instead.	0.2
attributes	object	Additional attributes needed to read and write the graph data.	null
detect_gzip	boolean	Enable or disable automatic gzip compression detection when loading graphs.	true
edge_destination_column	value	Name or index (starting from 1) of column corresponding to edge destination (for CSV format only).	null
edge_id_column	value	Name or index (starting from 1) of column corresponding to edge id (for CSV format only).	null
edge_id_strategy	enum[no_ids, keys_as_ids, unstable_generated_ids]	Indicates what ID strategy should be used for the edges of this graph. If not specified (or set to null), the strategy will be determined during loading or using a default value.	null
edge_id_type	enum[long]	Type of the edge ID. For homogeneous graphs, if not specified (or set to null), it will default to long.	null
edge_label_column	value	Name or index (starting from 1) of column corresponding to edge label (for CSV format only).	null
edge_props	array of object	Specification of edge properties associated with graph.	[]
edge_source_column	value	Name or index (starting from 1) of column corresponding to edge source (for CSV format only).	null
error_handling	object	Error handling configuration.	null
external_stores	array of object	Specification of the external stores where external string properties reside.	[]
format	enum[pgb, edge_list, adj_list, graphml, pg, rdf, two_tables]	Graph format to be used.	null
header	boolean	First line of file is meant for headers. For example, 'EdgeId, SourceId, DestId, EdgeProp1, EdgeProp2'.	false
keystore_alias	string	Alias to the keystore to use when connecting to database.	null
loading	object	Loading-specific configuration.	null

**Table 15-2 (Cont.) CSV Specific Options**

Field	Type	Description	Default
local_date_formats	array of string	array of local_date formats to use when loading and storing local_date properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string	[]
optimized_for	enum[read, updates]	Indicates if the graph should use data-structures optimized for read-intensive scenarios or for fast updates.	read
partition_while_loading	enum[by_label, no]	Indicates if the graph should be partitioned while loading.	null
password	string	Password to use when connecting to database.	null
point2d	string	Longitude and latitude as floating point values separated by a space.	0.0 0.0
separator	string	a series of single-character separators for tokenizing. The characters ", {, } and \n cannot be used as separators. Default value is ", " for CSV files, and "\t " for other formats. The first character will be used as a separator when storing.	null
storing	object	Storing-specific configuration.	null
time_format	array of string	The time format to use when loading and storing time properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string	[]
time_with_timezone_format	array of string	The time with timezone format to use when loading and storing time with timezone properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string.	[]
timestamp_format	array of string	The timestamp format to use when loading and storing timestamp properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string.	[]
timestamp_with_timezone_format	array of string	The timestamp with timezone format to use when loading and storing timestamp with timezone properties. See <a href="#">DateTimeFormatter</a> for documentation of the format string.	[]
vector_component_delimiter	character	Delimiter for the different components of vector properties.	;
vertex_id_column	value	Name or index (starting from 1) of column corresponding to vertex id (for CSV format only).	null
vertex_id_strategy	enum[no_ids, keys_as_ids, unstable_generated_ids]	Indicates what ID strategy should be used for the vertices of this graph. If not specified (or set to null), the strategy will be automatically detected.	null
vertex_id_type	enum[int, integer, long, string]	Type of the vertex ID. For homogeneous graphs, if not specified (or set to null), it will default to a specific value (depending on the origin of the data).	null



**Table 15-2 (Cont.) CSV Specific Options**

Field	Type	Description	Default
<code>vertex_labels_column</code>	value	Name or index (starting from 1) of column corresponding to vertex labels (for CSV format only).	<code>null</code>
<code>vertex_props</code>	array of object	Specification of vertex properties associated with graph.	<code>[]</code>

## 15.1.2 Specifying the File Path

The following examples show how to specify the file path for various file formats.

For formats that contain vertices and edges specified in one file (for example, `EdgeList`), use `uris` as shown in the following code:

```
{"uris":["path/to/file.format"]}
```

For formats that require separate files for edges and vertices (for example, `FlatFile`), use `vertex_uris` and `edge_uris` as shown in the following code:

```
{"vertex_uris":["vertices1.format","vertices2.format"],"edge_uris":
["edges1.format","edges2.format"]}
```

PGX will parse graphs in most of the plain text formats in parallel if the graph data is split into multiple files, as shown in the following code:

```
{"uris":["file1.format","file2.format",...,"fileN.format"]}
```

## 15.1.3 Supported File Access Protocols

The graph server (PGX) supports loading from graph configuration files and graph data files over various protocols and virtual file systems. The type of file system or protocol is determined by the scheme of the uniform resource identifier (URI):

- local file system (`file:`) - this is also the default if the given URI does not contain any scheme
- classpath (`classpath:` or `res:`)
- HDFS (`hdfs:`)
- HTTPS (`https:`)
- FTPS (`ftps:`)
- various archive formats (`zip:`, `jar:`, `tar:`, `tgz:`, `tbz2:`, `gz:` and `bz2:`). The URI format is `scheme://arch-file-uri[!absolute-path]` (if you would like to use the `!` as a literal file-name character it must be escaped using `%21`).  
For example, `jar:../lib/classes.jar!/META-INF/graph.json`.

Paths may be nested as in `tar:gz:https://anyhost/dir/mytar.tar.gz!/mytar.tar!/path/in/tar/graph.data`.

**Note:**

Relative paths are always resolved relative to the parent directory of the configuration file.

## 15.1.4 Plain Text Formats

The graph server (PGX) supports the following plain-text formats:

- Comma-Separated Values (CSV)
- Adjacency List (ADJ\_LIST)
- Edge List (EDGE\_LIST)
- Two Tables (TWO\_TABLES)
- Flat File (FLAT\_FILE)

### Parsing of Vertices

PGX supports three types of vertex identifies (`id`): `integer`, `long` and `string`. The type defaults to `integer`, but can be configured through the `vertex_id_type` option in the graph configuration.

### Parsing of Edges

Of the various formats and protocols supported by graph server (PGX), only CSV and flat file parsing support edge identifiers. For all other data sources, the `id` of an edge is PGX's internal `id`, which is an integer from zero to `num_edges - 1`.

### Parsing of Properties

`string` properties, spatial properties (currently only `point2d`) and temporal properties (`date`, `local_date`, `time`, `timestamp`, `time_with_timezone` and `timestamp_with_timezone`) must be quoted ("`<string>`") only if they contain a separator character (usually `,` for CSV and `' '` for Edge List and Adjacency List) or if they contain `"` or `\n`.

`date` properties are parsed using Java's `SimpleDateFormat` utility, instantiated with the format string `yyyy-MM-dd HH:mm:ss` unless specified otherwise in the graph configuration. All other types of temporal properties are parsed using Java's `DateTimeFormatter` utility.

`point2d` can be specified by its longitude followed by its latitude, separated by a space. Both longitude and latitude are doubles. For example, `"-74.0445 40.6892"` is the representation of a `point2d` instance representing the location of the Statue of Liberty.

Boolean values are interpreted as `true` if the value is `true` (ignoring case), `Y` (ignoring case) or `1`, `false` otherwise. The suggested notation for `false` is `false` (ignoring case), `N` (ignoring case) or `0`. All other types are parsed using the `parseXXX()` functions of its corresponding Java type, for example, `Integer.parseInt(...)` for integer types.

Vector properties are supported in the Adjacency List (ADJ\_LIST), Comma-Separated Values (CSV), Edge List (EDGE\_LIST), and Two Tables text (TWO\_TABLES) formats.

Vector properties with vector components of type `integer`, `long`, `float` and `double` can be loaded from these formats. In order to specify that a vertex or edge property is a vector property, the `dimension` field of the graph property configuration must be set to the dimension of the vector and be a strictly positive integer value. A vector value is represented in the supported text formats by the list of the vector components values separated by the vector component delimiter. By default the vector component delimiter is `;`, but this delimiter can be changed by changing the `vector_component_delimiter` graph configuration entry. Therefore a 3-dimensional vector of doubles could for example look like `0.1;0.0004;3.14` in the text file if the vector component delimiter is `;`.

## Separators

When using single file formats, IDs and properties are separated with `tab` or one single space (`"\t "`) by default, for multiple file formats comma (`" , "`) is used instead. However, PGX allows to configure the separator string.

## Parallel Loading

The following formats support parallel loading from multiple files:

- CSV (specify multiple files in `vertex_uris` and/or `edge_uris`)
- Adjacency List (specify multiple files in `uris`)
- Edge List (specify multiple files in `uris`)
- Two Tables (specify multiple files in `vertex_uris` and/or `edge_uris`)
- Flat File (specify multiple files in `vertex_uris` and/or `edge_uris`)

## Legend

The following abbreviations are used to specify text formats:

- V = Vertex Key
- VG = Neighbor Vertex
- VL = Vertex Labels
- VP = Vertex Property
- VPK = Vertex Property Key
- VPT = Vertex Property Type
- EL = Edge Label
- EP = Edge Property
- EPK = Edge Property Key
- EPT = Edge Property Type

For example `<V-2, VG-4>` or `<V-2, VG-4>` denotes the 4th neighbor of the 2nd vertex.

- [Comma-Separated Values \(CSV\)](#)
- [Adjacency List \(ADJ\\_LIST\)](#)
- [Edge List \(EDGE\\_LIST\)](#)
- [Two Tables \(TWO\\_TABLES\)](#)

### 15.1.4.1 Comma-Separated Values (CSV)

The CSV format is a text file format with vertices and edges stored in different files. Each line of the files represents a vertex or an edge. The vertex key and labels, the edge key, source, destination and label, and the attached properties are stored in the order specified by the file header (first line) and the configuration.

A graph with  $V$  vertices, having  $N$  vertex properties and  $K$  neighbors each, and  $E$  edges, having  $M$  edge properties, would be represented in CSV as shown:

vertices.csv

```
<V-1>,<VL-1>,<V-1, NP-1>,...,<V-1, NP-N>
<V-2>,<VL-2>,<V-2, NP-1>,...,<V-2, NP-N>
...
<V-V>,<VL-N>,<V-V, NP-1>,...,<V-V, NP-N>
```

edges.csv

```
<E-1>,<V-1>,<V-1, VG-1>,<EL-1>,<E-1, EP-1>,...,<E-1, EP-M>
...
<E-K>,<V-1>,<V-1, VG-K>,<EL-N>,<E-K, EP-1>,...,<E-K, EP-M>
<E-K+1>,<V-2>,<V-2, VG-1>,<EL-N+1>,<E-K+1, EP-1>,...,<E-K+1, EP-M>
...
<E-V*K>,<V-V>,<V-V, VG-K>,<EL-V*K>,<E-V*K, EP-1>,...,<E-V*K, EP-M>
```

#### Example 15-1 Loading graph from a CSV file with header details

The following examples shows a graph configuration file for loading a graph with two vertices and two edges:

vertices.csv

```
key,integer_prop,string_prop
1,33,"Alice"
2,42,"Bob"
```

edges.csv

```
source,dest,integer_prop,string_prop
1,2,0,"baz"
2,2,-12,"bat"
```

The corresponding graph configuration file is as shown:

```
{
 "format": "csv",
 "header": true,
 "vertex_id_column": "key",
 "edge_source_column": "source",
 "edge_destination_column": "dest",
 "vertex_uris": ["vertices.csv"],
 "edge_uris": ["edges.csv"],
 "vertex_props": [
 {
```

```

 "name": "integer_prop",
 "type": "integer"
 },
 {
 "name": "string_prop",
 "type": "string"
 }
],
"edge_props": [
 {
 "name": "integer_prop",
 "type": "integer"
 },
 {
 "name": "string_prop",
 "type": "string"
 }
]
}

```

### Example 15-2 Loading graph from a CSV file without header details

The following examples shows a graph configuration file for loading a graph with two vertices and two edges:

vertices.csv

```

1,33,"Alice"
2,42,"Bob"

```

edges.csv

```

1,2,0,"baz"
2,2,-12,"bat"

```

The corresponding graph configuration file is as shown:



#### Note:

The column indices are given in place of the column names.

```

{
 "format": "csv",
 "header": false,
 "vertex_id_column": 1,
 "edge_source_column": 1,
 "edge_destination_column": 2,
 "vertex_uris": ["vertices.csv"],
 "edge_uris": ["edges.csv"],
 "vertex_props": [
 {

```

```

 "name": "integer_prop",
 "type": "integer",
 "column": 2
 },
 {
 "name": "string_prop",
 "type": "string",
 "column": 3
 }
],
"edge_props": [
 {
 "name": "integer_prop",
 "type": "integer",
 "column": 3
 },
 {
 "name": "string_prop",
 "type": "string",
 "column": 4
 }
]
}

```

If no column indices are set in the configuration file, the columns are assumed to be in the following order:

- For vertex files: - Vertex ID - Vertex labels (if present) - Vertex properties in the order they are declared in the configuration
- For edge files: - Edge ID (if present) - Edge source - Edge destination - Edge label (if present) - Edge properties in the order they are declared in the configuration

Therefore the earlier configuration is equivalent to:

```

{
 "format": "csv",
 "header": false,
 "vertex_uris": ["vertices.csv"],
 "edge_uris": ["edges.csv"],
 "vertex_props": [
 {
 "name": "integer_prop",
 "type": "integer"
 },
 {
 "name": "string_prop",
 "type": "string"
 }
],
 "edge_props": [
 {
 "name": "integer_prop",
 "type": "integer"
 },
 {

```

```

 "name": "string_prop",
 "type": "string"
 }
}
}

```

### 15.1.4.2 Adjacency List (ADJ\_LIST)

The Adjacency List format is a text file format containing a list of neighbors from a vertex, per line. The format is extended to encode properties. The following shows a graph with  $V$  vertices, having  $N$  vertex properties and  $M$  edge properties:

```

<V-1> <V-1, VP-1> ... <V-1, VP-N> <V-1, VG-1> <EP-1> ... <EP-M> <V-1, VG-2> <EP-1> ...
<EP-M>
<V-2> <V-2, VP-1> ... <V-2, VP-N> <V-2, VG-1> <EP-1> ... <EP-M> <V-2, VG-2> <EP-1> ...
<EP-M>
...
<V-V> <V-V, VP-1> ... <V-V, VP-N> <V-V, VG-1> <EP-1> ... <EP-M> <V-V, VG-2> <EP-1> ...
<EP-M>

```

#### Note:

Trailing separators will be considered as errors. For example, if whitespace is used to separate the properties, any trailing whitespace will cause an exception to be raised.

#### Example 15-3 Graph in Adjacency List Format

This example shows a graph with 4 vertices (1, 2, 3 and 4), each having a double and a string property, and 3 edges, each having a boolean and a date property, encoded in Adjacency List format:

```

1 8.0 "foo"
2 4.3 "bar" 1 false "1985-10-18 10:00:00"
3 6.1 "bax" 2 true "1961-12-30 14:45:14" 4 false "2001-01-15 07:00:43"
4 17.78 "f00"

```

#### Note:

ADJ\_LIST is more space efficient than EDGE\_LIST. This is because vertices are first defined and then the edges are being created, indicating that we are repeating each vertex at least once.

### 15.1.4.3 Edge List (EDGE\_LIST)

The Edge List format is a text file format starting with a section with one vertex per line, followed by a section with one edge per line. If a vertex does not have any labels or properties, it is possible to omit the vertex in the first section, but still specify edges for the vertex in the second section.

```

EdgeList := {Vertex '\n'}* '\n' {Edge '\n'}*

Vertex := VertexId '*' VertexLabels? PropertyValue*
VertexId := Integer | Long | String
VertexLabels := {' String* '}

Edge := SrcVertex DstVertex EdgeLabel? PropertyValue*
SrcVertex := VertexId
DstVertex := VertexId
EdgeLabel := String

PropertyValue := Integer | Long | Double | Float | Boolean | String | Date

```

The vertices start with an identifier (`VertexId`), followed by a `*`, an optional set of vertex labels (`VertexLabels?`) and the vertex properties (`PropertyValue*`). A vertex identifier is either an `Integer`, a `Long`, or a `String`. Furthermore, vertex labels are zero or more `Strings` between curly braces (`{' String* '}`).

The edges start with source and destination vertex identifiers (`SrcVertex DstVertex`), followed by optional edge label (`EdgeLabel?`) and the edge properties (`PropertyValue*`). The edge label is a `String`.

#### Example 15-4 Graph in Edge List format

This example shows a graph with two vertices and two edges, with labels and properties:

```

1 * { "Person" "Male" } "Mario" 15
2 * { "Person" "Male" } "Luigi" 14
1 2 "likes" 3.5
2 1 "likes" 2.1

```

The two vertices (lines 1-2) have identifiers 1 and 2 and both have the labels "Person" and "Male", a string property ("Mario" and "Luigi") and an integer property (15 and 14). There is an edge from vertex 1 to vertex 2 (line 3) with label "likes" and a double property with value 3.5, and another edge from vertex 2 to vertex 1 with label "likes" and a double property with value 2.1.

The following shows the corresponding graph configuration:

```

{
 "format":"edge_list",
 "uri":"example.edgelist",
 "vertex_id_type":"long",
 "vertex_labels":true,
 "edge_label":true,
 "vertex_props":[
 {
 "name":"name",
 "type":"string"
 },
 {
 "name":"age",
 "type":"int"
 }
]
},

```



```

"edge_props":[
 {
 "name":"rating",
 "type":"double"
 }
],
"loading_options": {
 "load_vertex_labels":true,
 "load_edge_label":true
},
"separator":" "
}

```

#### 15.1.4.4 Two Tables (TWO\_TABLES)

When configured to use `file` as datastore, the Two Tables format becomes a text file format similar to the Edge List format, with the only difference that the vertices and edges are stored in two different files. The vertices file contains vertex IDs followed by vertex properties. The edges file contains the source vertices and target vertices, followed by edge properties.

A graph with  $V$  vertices, having  $N$  vertex properties and  $M$  edge properties would be represented in two files as shown in the following:

```

vertices.ttt:

<V-1> <V-1, NP-1> ... <V-1, NP-N>
<V-2> <V-2, NP-1> ... <V-2, NP-N>
...
<V-V> <V-V, NP-1> ... <V-V, NP-N>

edges.ttt:

<V-1> <V-1, VG-1> <EP-1> ... <EP-M>
<V-1> <V-1, VG-2> <EP-1> ... <EP-M>
...
<V-V> <V-V, VG-1> <EP-1> ... <EP-M>

```

#### Example 15-5 Graph in Two Tables Text format

The following example shows the graph of 4 vertices (1, 2, 3 and 4), each having a `double` and a `string` property, and 3 edges, each having a `boolean` and a `date` property, encoded in Two Tables Text format:

```

vertices.ttt:

1 8.0 "foo"
2 4.3 "bar"
3 6.1 "bax"
4 17.78 "f00"

edges.ttt:

2 1 false "1985-10-18 10:00:00"
3 2 true "1961-12-30 14:45:14"
3 4 false "2001-01-15 07:00:43"

```

 **Note:**

If you are planning on storing big graphs you must consider Two Tables Text format in order to save disk space.

## 15.1.5 XML File Formats

### Graph ML

The graph server (PGX) supports loading graphs from files using the XML-based Graph ML format. Graphs already in memory may also be exported into GraphML files. See [GraphML specification](#) for a detailed description of the XML schema.

### PGX GraphML Limitation

PGX does not support all features of the GraphML format. Some of the limitations are:

- If the graph is undirected (`edgedefault="undirected"`), then edge properties are not supported
- All vertices (edges) must have the same amount and type of vertex (edge) properties
- `port`, `default`, and `hyperedge` are not supported

### Example 15-6

The following example graph consists of 3 vertices and 3 edges. Each vertex has an integer property named `number` and each edge has a string property named `label`. Note that the edges are directed and that the strings for the property do not have to be put in (double) quotation marks.

```
<?xml version="1.0" encoding="UTF-8"?>
<graphml xmlns="http://graphml.graphdrawing.org/xmlns"
 <key attr.name="number" attr.type="integer" for="node"
id="number"/>
 <key attr.name="label" attr.type="string" for="edge" id="label"/>
 <graph edgedefault="directed">
 <node id="1">
 <data key="number">2</data>
 </node>
 <node id="2">
 <data key="number">45</data>
 </node>
 <node id="3">
 <data key="number">83</data>
 </node>
 <edge target="2" source="1">
 <data key="label">this graph</data>
 </edge>
 <edge source="3" target="2">
 <data key="label">forms a</data>
 </edge>
 <edge target="1" source="3">
 <data key="label">triangle</data>
 </edge>
 </graph>
</graphml>
```

```

 </edge>
 </graph>
</graphml>

```

### ▲ Caution:

Due to the verbose nature of XML, the GraphML format comes with a large overhead compared to other file-based graph formats. You must use a different format if you want to consider the load or store performance and file size as important factors.

## 15.1.6 Binary File Formats

### PGX Binary Format (PGB)

PGX binary format (.pgb) is the proprietary binary format for graph server (PGX), which allows fast and efficient file processing. Fundamentally, the file is a binary dump of the graph and property data. Bytes are written in network byte order (big endian).

### Type Encoding

**Table 15-3** Type Encoding

Value	Type	Size in bytes
0	Boolean	1
1	Integer	4
2	Long	8
3	Float	4
4	Double	8
7	String	varies
11	Vertex labels	varies
13	Local date	4
14	Time	4
15	Timestamp	8
16	Time with time zone	8
17	Timestamp with time zone	12
18	Vector property	variable: <sizeof component-type> * <dimension>

### File Layout

**Table 15-4** File Layout

Size in bytes	Description	Required	Comment
4	magic word	Yes	0x99191191

**Table 15-4 (Cont.) File Layout**

Size in bytes	Description	Required	Comment
4	vertex size	Yes	Allowed values are 4 and 8.
4	edge size	Yes	Allowed values are 4 and 8.
<vertex size>	number of vertices	Yes	
<edge size>	number of edges	Yes	
<edge size> * (<numVertices> + 1)	edge begin array	Yes	
<vertex size> * <numEdges>	destination vertex array	Yes	
1	component bitmap	Yes	<ul style="list-style-type: none"> <li>• 0x0001: node keys</li> <li>• 0x0002: vertex labels</li> <li>• 0x0004: edge label</li> <li>• 0x0008: edge keys</li> <li>• other bits: reserved</li> </ul>
4	vertexKey type	No	Only present if <i>component bitmap</i> & 0x0001 == 0x0001. See <a href="#">Table 15-3</a> for type encoding.
<vertex key layout>	vertex keys	No	Only present if <i>component bitmap</i> & 0x0001 == 0x0001.
4	edgeKey type	No	Only present if <i>component bitmap</i> & 0x0008 == 0x0008. See <a href="#">Table 15-3</a> for type encoding
<numEdges> * 8	edge keys	No	Only present if <i>component bitmap</i> & 0x0008 == 0x0008.
4	number of vertex properties	Yes	
<num vertex properties> * <property layout>	property data	Yes	See <a href="#">Table 15-10</a> .
4	number of edge properties	Yes	
<num edge properties> * <property layout>	property data	Y	See <a href="#">Edge Property Layout</a> .
<vertex labels layout>	vertex labels	No	Only present if <i>component bit</i> & 0x0002 == 0x0002.
<edge labels layout>	edge label	No	Only present if <i>component bit</i> & 0x0004 == 0x0004.
4	number of shared pools	Yes	
<shared pools size>	shared pools	No	
<property names size>	property names	No	Only present if <i>component bit</i> & 0x0010 == 0x0010. See <a href="#">Table 15-19</a> .

### Vertex Key Layout

The layout of vertex keys depends on the vertexKey type. PGB supports integer, long and string vertex keys.

**Table 15-5 Integer Vertex Keys**

Size in bytes	Description	Required	Comment
<numVertices> * 4	key data	Yes	For each vertex, the corresponding integer key value.

**Table 15-6 Long Vertex Keys**

Size in bytes	Description	Required	Comment
<numVertices> * 8	key data	Yes	For each vertex, the corresponding long key value.

**Table 15-7 String Vertex Keys**

Size in bytes	Description	Required	Comment
4	compression scheme	Yes	reserved (must be 0)
8	property size	Yes	size of each element in bytes in the following data
<number of keys> * <string key element layout>	string key data	Yes	content of the vertex keys (see <a href="#">Table 15-5</a> )

**Table 15-8 String Key Element Layout**

Size in bytes	Description	Required	Comment
4	string length	Yes	length of the string in bytes
<string length>	string key data	Yes	content of the string as bytes, <b>No zero-character</b>

### Property Layout

The following shows the special layout for string properties, and for vector properties:

**Table 15-9 Primitive Type Layout**

Size in bytes	Description	Required	Comment
4	property type	Yes	See <a href="#">Table 15-3</a> for type encoding.
8	property size	Yes	Size of the property data in bytes
<property size>	property data	Yes	Stored as <numVertices/numEdges> * <type size>

**Table 15-10 Vector Property Layout**

Size in bytes	Description	Comment
4	vector type mark	Always equal to 18.
8	size of vector property data and extra fields	$dataSize = \langle sizeof\ component\text{-}type \rangle * \langle dimension \rangle + 8$ (The 8 extra bytes are for the added following 2 extra fields in the vector property header.)
4	vector component data type	Valid types are <code>integer</code> , <code>long</code> , <code>float</code> , <code>double</code> . Encoded with the value specified in <a href="#">Table 15-3</a> .
4	vector dimension	Number of components per vector value. Must be greater than 0 to be a valid vector property.
$dataSize - 8$	data	Stored as array of length $i * \backslash$ in which the value of the $j$ -th component of the vector for the $i$ -th entity is at position $i * \backslash + j$ .

**Table 15-11 String Type Layout**

Size in bytes	Description	Required	Comment
4	property type	Yes	Must be 7.
8	property size	Yes	Size of the following data in bytes.
1	reserved	Yes	Reserved (must be 0).
$\langle dictionary\ layout \rangle$	dictionary	Yes	String dictionary used in the property
$\langle numVertices / numEdges \rangle * 8$	property content	Yes	Content of the string property, stored as IDs that refer to the strings in the dictionary.

**Table 15-12 String Dictionary Layout**

Size in bytes	Description	Required	Comment
1	reserved	Yes	Reserved (must be 0).
8	number of strings	Yes	Number of strings in the following dictionary.
$\langle number\ of\ strings \rangle * \langle dictionary\ element\ layout \rangle$	dictionary data	Yes	See <a href="#">Table 15-13</a> .

**Table 15-13 String Dictionary Element Layout**

Size in bytes	Description	Required	Comment
8	string id	Yes	Unique ID of the string.

**Table 15-13 (Cont.) String Dictionary Element Layout**

Size in bytes	Description	Required	Comment
4	string length	Yes	Length of the string in bytes.
<string length>	string data	Yes	Content of the string as bytes, <b>No zero-character</b>

**Vertex Labels Layout****Table 15-14 Vertex Labels Layout**

Size in bytes	Description	Required	Comment
4	type	Yes	Must be 11.
8	size	Yes	Size of the following data in bytes.
<dictionary layout>	dictionary	Yes	String dictionary used in the vertex labels.
<numVertices + 1> * 8	string id begin array	Yes	<string ids> offset array for each vertex.
8	number of string ids	Yes	The number of string ids.
<number of string ids> * 8	string ids	Yes	Array of string ids in the string dictionary.

**Edge Label Layout**

The edge label layout follows the string type layout.

**Shared Pools Layout****Table 15-15 Shared Pools Layout**

Size in bytes	Description	Required	Comment
1	type	Yes	1: enum, 2: prefixed

**Table 15-16 Type == Enum**

Size in bytes	Description	Required	Comment
8	num strings	Yes	
<number of strings> * <string table layout>	dictionary data	Yes	See <a href="#">Table 15-18</a> .

**Table 15-17 Type == Prefix**

Size in bytes	Description	Required	Comment
8	num prefixes	Yes	
<number of prefixes> *	dictionary data	Yes	See <a href="#">Table 15-18</a> .
<string table layout>			
8	num suffixes	Yes	
<number of suffixes> *	dictionary data	Yes	See <a href="#">Table 15-18</a> .
<string table layout>			

**Table 15-18 String Table for Shared Pools**

Size in bytes	Description	Required	Comment
8	string id	Yes	String can be literal (in case of enum) or prefix/suffix (in case of prefix).
4	string length	Yes	
<string length>	string data	Yes	

### Property Names Layout

**Table 15-19 Property Names Layout**

Size in bytes	Description	Required	Comment
8	size	Yes	String can be literal (in case of enum) or prefix/suffix (in case of prefix).
<sum of size of vertex property names>	vertex property names	No	Follows the String Key Element Layout. See <a href="#">Table 15-8</a> .
<sum of size of edge property names>	edge property names	No	Follows the String Key Element Layout. See <a href="#">Table 15-8</a> .

## 15.2 Loading Graph Data in Parallel from Multiple Files

You can load a graph in parallel using multiple files.

The following example demonstrates how to load graph data from multiple files.



For example, consider a vertex file split into four partitions as shown:

```
vertex_file1

1,Color,1,red,,
2,Color,1,yellow,,

vertex_file2

3,Color,1,blue,,
4,Color,1,green,,

vertex_file3

5,Color,1,orange,,
6,Color,1,white,,

vertex_file4

7,Color,1,black,,
```

The edge file is split into two partitions as shown:

```
edge_file1

1,1,2,edge1,Weight,4,,1.0,
2,2,3,edge2,Weight,4,,2.0,
3,3,4,edge3,Weight,4,,3.0,

edge_file2

4,4,5,edge4,Weight,4,,4.0,
5,5,6,edge5,Weight,4,,5.0,
6,6,7,edge6,Weight,4,,6.0,
```

The following graph configuration can be used to load the graph data from four vertex files and two edge files into the same graph. Note that all the `uris` are specified inside the JSON graph configuration.

```
{
 "format": "flat_file",
 "vertex_uris": ["vertex_file1", "vertex_file2", "vertex_file3",
"vertex_file4"],
 "edge_uris": ["edge_file1", "edge_file2"],
 "separator": ",",
 "edge_props": [
 {
 "name": "Weight",
 "type": "double"
 }
]
}
```

```
 }
],
 "vertex_props": [
 {
 "name": "Color",
 "type": "string"
 }
]
}
```

You can also create a graph configuration with multiple file partitions using Java as shown:

```
FileGraphConfig config = GraphConfigBuilder
 .forFileFormat(Format.FLAT_FILE)
 .setSeparator(",")
 .addVertexUri("vertex_file1")
 .addVertexUri("vertex_file2")
 .addVertexUri("vertex_file3")
 .addVertexUri("vertex_file4")
 .addEdgeUri("edge_file1")
 .addEdgeUri("edge_file2")
 .addVertexProperty("Color", PropertyType.STRING)
 .addEdgeProperty("Weight", PropertyType.DOUBLE)
 .build();
```

 **Note:**

The graph configuration in the preceding codes include one double edge property named "Weight" and one string vertex property named "Color".

You can now load the graph data from the files as explained in [Creating a graph using graph builder API](#).

The graph server (PGX) will automatically load the graph in parallel, using one thread for each file. This means that a graph can be loaded in parallel with as many threads as files are given depending on the configured parallelism for the graph server (PGX) instance.

 **Note:**

Since the graph config will be used for all of the specified files, it is crucial to use the same format for all these files, that is, using the same separator, having the same defined properties, complying with the same format specification.

## 15.3 Exporting Graphs Into a File

The graph server (PGX) allows the client to export a currently loaded graph into a file.

Using the `store()` method on any `PgxGraph` object, the client can specify which file format to store the graph in. The client can also dynamically select the set of properties to be stored with the graph, that is, not all the properties need to be exported. The client can specify a `CompressionScheme` to use when storing as shown:

**Table 15-20 Files CompressionScheme**

CompressionScheme	Supported Formats
NONE	All formats
GZIP	ADJ_LIST, EDGE_LIST, FLAT_FILE, TWO_TABLES (text)

The client can export to multiple files as well.

When PGX exports the specified graph into a file, PGX also creates a graph config which the client receives as return value. This is to help loading the created graph instance later.

When exporting graph data into multiple files a `FileGraphStoringConfig` can be used which contains the following JSON fields:

**Table 15-21 Graph Configuration when Exporting Graph into Multiple Files**

Field	Type	Description	Default
<code>base_path</code>	<code>string</code>	Base path to use for storing a graph; file paths will be constructed using the following format <code>_. _</code> , that is, <code>parent_path/my_graph_1.edges</code> .	<code>null</code>
<code>compression_scheme</code>	<code>enum[none, gzip]</code>	The scheme to use for compression, or none to disable compression.	<code>none</code>
<code>delimiter</code>	<code>character</code>	Delimiter character used as separator when storing. The characters <code>"</code> , <code>{</code> , <code>}</code> and <code>\n</code> cannot be used as delimiters.	<code>null</code>
<code>edge_extension</code>	<code>string</code>	The extension to use when creating edge file partitions.	<code>edges</code>
<code>initial_partition_index</code>	<code>integer</code>	The value used as initial partition index, that is, <code>initial_partition_index=1024 -&gt; my_graph_1024.edges</code> , <code>my_graph_1025.edges</code> , ...	<code>1</code>
<code>num_partitions</code>	<code>integer</code>	The number of partitions that should be created, when exporting to multiple files.	<code>1</code>

**Table 15-21 (Cont.) Graph Configuration when Exporting Graph into Multiple Files**

Field	Type	Description	Default
row_extension	string	The extension to use when creating row file partitions.	rows
vertex_extension	string	The extension to use when creating vertex file partitions.	nodes

- [Exporting a Graph to Disk](#)

## 15.3.1 Exporting a Graph to Disk

You can save a graph loaded into memory to the disk in various formats. Therefore you can make sub-graphs and graph data computed at run time through analytics persistent, for future use. The resulting file can be used later as input for the graph server (PGX).

Consider the following example where a graph is loaded into memory and PageRank analysis is executed on the graph.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

### JShell

```
var g = session.readGraphWithProperties("<path_to_json>")
var rank = analyst.pagerank(g, 0.001, 0.85, 100)
```

### Java

```
PgxGraph g = session.readGraphWithProperties("<path_to_json>");
Analyst analyst = session.createAnalyst();
VertexProperty<Integer, Double> rank = analyst.pagerank(g, 0.001,
0.85, 100);
```

### Python

```
g = session.read_graph_with_properties("<path_to_json>")
analyst = session.create_analyst()
rank = analyst.pagerank(g, 0.001, 0.85, 100)
```

---

You can now store the graph, together with the result of the PageRank analysis and all original edge properties, as a file in edge-list format, on disk. When a graph is stored, you need to specify the graph format, a path where the file should be stored, the properties to store and a flag that specifies whether or not a file should be overwritten should a file with the same name already exist.

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
var config = g.store(Format.EDGE_LIST, "<file-path>", List.of(rank),
EdgeProperty.ALL, false)
```

## Java

```
var config = g.store(Format.EDGE_LIST, "<file-path>", List.of(rank),
EdgeProperty.ALL, false);
```

## Python

```
config = g.store('edge_list', "<file-path>", vertex_properties = [rank],
overwrite= False)
```

---

The graph data can now be found under the file path. The graph configuration returned by the `store` method can be used to load the new graph back into memory. To persist the graph configuration to disk as well, you can use the config's `toString` method to get a JSON representation:

- 
- [JShell](#)
  - [Java](#)
  - [Python](#)

## JShell

```
var path = Paths.get("<file-path>")
Files.writeString(path, config.toString())
```

## Java

```
import apache.commons.io.*; // PGX contains a version of Apache
Commons IO
...
FileUtils.write(new File("<file-path>"), config.toString());
```

## Python

```
with open("<file-path>", "w"):
 f.write(str(config))
```

---

# 15.4 Exporting a Graph into Multiple Files

You can store a graph into multiple files using the `store` method. Most parameters are the same, as if storing to a single file. However, the main difference lies in specifying how to partition the data.

You can partition the data in either of the following two ways:

- specifying a `FileGraphStoringConfig` (see [Table 15-21](#) for more information)
- specifying a base path and the number of partitions

### Export into Multiple Files Using `FileGraphStoringConfig`

You can specify a more detailed way of creating the multiple partitions used to store the graph by using the `FileGraphStoringConfig`. You can create a `FileGraphStoringConfig` object using a `FileGraphStoringConfigBuilder`.

For example, the following code specifies that the storing should be done into four partitions using the specified base path and using zero as the initial index for the partitioning. It also contains the file extension to use for vertex files and for edge files and finally it sets comma as the delimiter to be used when storing the graph data:

```
FileGraphStoringConfig storingConfig = new
FileGraphStoringConfigBuilder(basePath) //
 .setNumPartitions(4) //
 .setInitialPartitionIndex(0) //
 .setVertexExtension(vertexExtension) //
 .setEdgeExtension(edgeExtension) //
 .setDelimiter(',') //
 .build();
```

You can also partition all tables equally using the `numPartitions` parameter. This implies that all tables are exported into the same number of files.

If you do not want to partition the tables equally, you can either create one `PartitionedGraphConfig` which contains for each provider a `FileGraphStoringConfig` (see [Table 15-21](#)) or we can use a version of `store()` that

takes two maps of `FileGraphStoringConfigs`, one for the vertex tables and one for the edge tables.

For the first option, you can create for each vertex and edge table a `FileGraphStoringConfig` and put it into a `FileEntityProviderConfig` using `setStoringOptions` in the builder of `FileEntityProviderConfig`. The providers are then added to the `PartitionedGraphConfig` as edge and vertex providers using `addVertexProvider()` and `addEdgeProvider()` in the builder of `PartitionedGraphConfig`. Later you can use the `store()` method which takes the `PartitionedGraphConfig` as parameter.

The second option creates for every edge and vertex table a storing configuration, adds those into a vertex provider and an edge provider map and calls the corresponding `store()` method with these maps as parameters.

For example:

```
FileGraphStoringConfig vertexStoringConfig1 = new
FileGraphStoringConfigBuilder(basePath + "_vertexTable1") //
 .setNumPartitions(4) //
 .setInitialPartitionIndex(0) //
 .setVertexExtension(vertexExtension) //
 .setDelimiter(',') //
 .build();

FileGraphStoringConfig vertexStoringConfig2 = new
FileGraphStoringConfigBuilder(basePath + "_vertexTable2") //
 .setNumPartitions(4) //
 .setInitialPartitionIndex(0) //
 .setVertexExtension(vertexExtension) //
 .setDelimiter(',') //
 .build();

FileGraphStoringConfig edgeStoringConfig1 = new
FileGraphStoringConfigBuilder(basePath + "_edgeTable1") //
 .setNumPartitions(4) //
 .setInitialPartitionIndex(0) //
 .setEdgeExtension(edgeExtension) //
 .setDelimiter(',') //
 .build();

Map<String, FileGraphStoringConfig> vertexStoringConfigs = new HashMap<>();
vertexStoringConfigs.put("vertexTable1", vertexStoringConfig1);
vertexStoringConfigs.put("vertexTable2", vertexStoringConfig2);

Map<String, FileGraphStoringConfig> edgeStoringConfigs = new HashMap<>();
edgeStoringConfigs.put("edgeTable1", edgeStoringConfig);
```

### Export into Multiple Files without `FileGraphStoringConfig`

If you only need to specify how many partitions are required and the base name to be used, it is simpler to use `store()` method by only specifying those parameters. Following this procedure, the graph server (PGX) will use defaults for the other fields. See [Table 15-21](#) for more information on default values.

### Export into Multiple Files Using a Graph Configuration Object

An alternate way for exporting into multiple files is by creating a `FileGraphStoringConfig` and putting it into a `Graph Configuration` object using `setStoringOptions` in its builder, and then using the corresponding version of the `store()` method.



# 16

## Log Management in the Graph Server (PGX)

The graph server (PGX) internally uses the SLF4J interface with Log4j as the default logger implementation.

- [Configuring Log4j Logging](#)

### 16.1 Configuring Log4j Logging

The default log4j logging configuration file is located in `/etc/oracle/graph/log4j2-server.xml`. This configuration file contains the target location for the logs in `/var/log/oracle/graph/`. Additionally, the rolling file appenders are also defined in this configuration file.

#### Note:

- Log4j is configured to roll the log files based on both log size (250 MB) and date.
- Log files are automatically saved in a compressed format in subdirectories, one directory per month. There can be multiple files on a given day.
- Also, each startup of the graph server(PGX) triggers a new log file.

The log4j configuration file is picked up automatically by the the graph server(PGX). To use this configuration in your java application, you can set the `log4j.configurationFile` system variable when launching the JVM:

```
java -Dlog4j.configurationFile=$PGX_HOME/conf/log4j2.xml ...
```

#### Changing Logging Level During a JShell Session

When connected to the graph server using JShell, you can use the `loglevel(String loggerName, String levelName)` function to quickly change the logging level of any logger. For example:

```
loglevel("oracle.pgx", "debug")
loglevel("ROOT", "info")
loglevel("org.apache.hadoop", "off")
```

#### Changing Slf4j Implementation

You can replace the log4j JARs in `$PGX_HOME/third-party` with your own slf4j implementation. You must only place your JAR files in `$PGX_HOME/third-party` and it will get wild-card included when the graph shell client is started.

### Logging in a Web Application Server

The `graph-server-<version>-pgx<version>.war` file in the `oracle-graph-webapps-<version>.zip` download package contains the `log4j2.xml`. This file determines what should be logged in the web application running on the application server of your choice. The file is located in the folder `WEB-INF/classes` inside the `graph-server-<version>-pgx<version>.war` file. By default, only errors are logged. But you can change this file if you want more logging in your web server. You must restart the web server after you change the file, for the change to take effect.

# Part III

## Supplementary Information for Property Graph Support

This document has the following appendixes.

- [Mapping Graph Server Roles to Default Privileges](#)
- [Disabling Transport Layer Security \(TLS\) in Graph Server](#)

# A

## Mapping Graph Server Roles to Default Privileges

Installing the PL/SQL packages of the Oracle Graph Server and Client distribution on the target Oracle Database, automatically creates the following roles and assigns the default permissions as shown in the following table:

**Table A-1 Mapping Graph Server Roles to Default Privileges**

Roles	Description	Permission
GRAPH_ADMINISTRATOR	User who performs operations on the in-memory graph server (PGX) using the Java API. (As compared to running start and stop operations as an OS user.)	PGX_SESSION_CREATE PGX_SERVER_GET_INFO PGX_SERVER_MANAGE
GRAPH_DEVELOPER	User who creates graphs, publishes graphs, modifies graphs, queries graphs, and views graphs using the Java API or SQLcl or the graph visualization application.	PGX_SESSION_CREATE PGX_SESSION_NEW_GRAPH PGX_SESSION_GET_PUBLISHED_GRAPH PGX_SESSION_MODIFY_MODEL PGX_SESSION_READ_MODEL
GRAPH_USER	User who queries graphs and views graphs Java API or SQLcl or the graph visualization application.	PGX_SESSION_CREATE PGX_SESSION_GET_PUBLISHED_GRAPH

# B

## Disabling Transport Layer Security (TLS) in Graph Server

For demonstration or evaluation purposes, it is possible to turn off transport layer security (TLS) of the graph server.

### **Caution:**

This is **not** recommended for production. In a secure configuration, the server must always have TLS enabled.

The following instructions only apply if you installed the graph server via the RPM package.

### **Note:**

If you deployed the graph server into your own web server (e.g Weblogic or Apache Tomcat), please refer to the manual of your web server for TLS configuration.

1. Edit `/etc/oracle/graph/server.conf` to change `enable_tls` to `false`.
2. Edit the `WEB-INF/web.xml` file inside the WAR file in `/opt/oracle/graph/graphviz` and configure cookies to be sent over non-secure connections by setting `<secure>>false</secure>` as follows:

```
<session-config>
 <tracking-mode>COOKIE</tracking-mode>
 <cookie-config>
 <secure>>false</secure>
 </cookie-config>
 ...
</session-config>
```

3. Additionally, replace `https` with `http` in the `pgx.base_url` property in the same `WEB-INF/web.xml` file. For example:

```
<context-param>
 <param-name>pgx.base_url</param-name>
 <param-value>http://localhost:7007</param-value>
</context-param>
```

4. Restart the server.

```
sudo systemctl restart pgx
```

The graph server now accepts connections over HTTP instead of HTTPS.

On Oracle Linux 7, you can execute the following script to perform the preceding four steps all at once:

```
echo "$(jq '.enable_tls = false' /etc/oracle/graph/server.conf)"
> /etc/oracle/graph/server.conf
WAR=$(find /opt/oracle/graph/graphviz -name '*.war')
TMP=$(mktemp -d)
cd $TMP
unzip $WAR WEB-INF/web.xml
sed -i 's|<secure>>true</secure>|<secure>>false</secure>|' WEB-INF/
web.xml
sed -i 's|https://|http://|' WEB-INF/web.xml
sudo zip $WAR WEB-INF/web.xml
rm -r $TMP
sudo systemctl restart pgx
```

# Index

## A

---

ANALYZE\_PG procedure, [9-2](#)  
automatic delta refresh, [4-60](#)

## C

---

CF procedure, [9-4](#)  
CF\_CLEANUP procedure, [9-7](#)  
CF\_PREP procedure, [9-9](#)  
CLEAR\_PG procedure, [9-10](#)  
CLEAR\_PG\_INDICES procedure, [9-11](#)  
CLONE\_GRAPH procedure, [9-11](#)  
collaborative filtering, [9-4](#), [9-7](#), [9-9](#)  
connected components  
    finding, [9-32](#)  
COUNT\_TRIANGLE function, [9-12](#)  
COUNT\_TRIANGLE\_CLEANUP procedure, [9-13](#)  
COUNT\_TRIANGLE\_PREP procedure, [9-14](#)  
COUNT\_TRIANGLE\_RENUM function, [9-16](#)  
CREATE\_EDGES\_TEXT\_IDX procedure, [9-17](#)  
CREATE\_PG procedure, [9-18](#)  
CREATE\_PG\_SNAPSHOT\_TAB procedure, [9-19](#)  
CREATE\_PG\_TEXTIDX\_TAB procedure, [9-21](#)  
CREATE\_STAT\_TABLE procedure, [9-22](#)  
CREATE\_SUB\_GRAPH procedure, [9-23](#)  
CREATE\_VERTICES\_TEXT\_IDX procedure,  
    [9-24](#)

## D

---

DROP\_EDGES\_TEXT\_IDX procedure, [9-26](#)  
DROP\_PG procedure, [9-26](#)  
DROP\_PG\_VIEW procedure, [9-27](#)  
DROP\_VERTICES\_TEXT\_IDX procedure, [9-27](#)

## E

---

edge table statistics  
    exporting, [9-30](#)  
    importing, [9-51](#)  
ESTIMATE\_TRIANGLE\_RENUM function, [9-28](#)  
EXP\_EDGE\_TAB\_STATS procedure, [9-30](#)  
EXP\_VERTEX\_TAB\_STATS procedure, [9-31](#)

## F

---

FIND\_CC\_MAPPING\_BASED procedure, [9-32](#)  
FIND\_CLUSTERS\_CLEANUP procedure, [9-33](#)  
FIND\_CLUSTERS\_PREP procedure, [9-34](#)  
FIND\_SP procedure, [9-36](#)  
FIND\_SP\_CLEANUP procedure, [9-37](#)  
FIND\_SP\_PREP procedure, [9-38](#)

## G

---

geometries  
    getting, [9-39](#), [9-41](#)  
    getting from longitude and latitude, [9-44](#)  
    WKT, [9-48](#), [9-49](#)  
GET\_BUILD\_ID function, [9-39](#)  
GET\_GEOMETRY\_FROM\_V\_COL function,  
    [9-39](#)  
GET\_GEOMETRY\_FROM\_V\_T\_COLS function,  
    [9-41](#)  
GET\_LATLONG\_FROM\_V\_COL function, [9-42](#),  
    [9-45](#)  
GET\_LATLONG\_FROM\_V\_T\_COLS function,  
    [9-43](#)  
GET\_LONG\_LAT\_GEOMETRY function, [9-44](#)  
GET\_LONGLAT\_FROM\_V\_T\_COLS function,  
    [9-46](#)  
GET\_SCN function, [9-47](#)  
GET\_VERSION function, [9-47](#)  
GET\_WKTGEOMETRY\_FROM\_V\_COL function,  
    [9-48](#)  
GET\_WKTGEOMETRY\_FROM\_V\_T\_COLS  
    function, [9-49](#)  
GRANT\_ACCESS procedure, [9-50](#)

## I

---

IMP\_EDGE\_TAB\_STATS procedure, [9-51](#)  
IMP\_VERTEX\_TAB\_STATS procedure, [9-52](#)  
in-memory Graph server (PGX), [4-1](#)

## O

---

OPG\_APIS package

- ANALYZE\_PG, [9-2](#)
- CF, [9-4](#)
- CF\_CLEANUP, [9-7](#)
- CF\_PREP, [9-9](#)
- CLEAR\_PG, [9-10](#)
- CLEAR\_PG\_INDICES, [9-11](#)
- CLONE\_GRAPH, [9-11](#)
- COUNT\_TRIANGLE, [9-12](#)
- COUNT\_TRIANGLE\_CLEANUP, [9-13](#)
- COUNT\_TRIANGLE\_PREP, [9-14](#)
- COUNT\_TRIANGLE\_RENUM, [9-16](#)
- CREATE\_EDGES\_TEXT\_IDX, [9-17](#)
- CREATE\_PG, [9-18](#)
- CREATE\_PG\_SNAPSHOT\_TAB, [9-19](#)
- CREATE\_PG\_TEXTIDX\_TAB, [9-21](#)
- CREATE\_STAT\_TABLE, [9-22](#)
- CREATE\_SUB\_GRAPH, [9-23](#)
- CREATE\_VERTICES\_TEXT\_IDX, [9-24](#)
- DROP\_EDGES\_TEXT\_IDX, [9-26](#)
- DROP\_PG, [9-26](#)
- DROP\_PG\_VIEW, [9-27](#)
- DROP\_VERTICES\_TEXT\_IDX, [9-27](#)
- ESTIMATE\_TRIANGLE\_RENUM, [9-28](#)
- EXP\_EDGE\_TAB\_STATS, [9-30](#)
- EXP\_VERTEX\_TAB\_STATS, [9-31](#)
- FIND\_CC\_MAPPING\_BASED, [9-32](#)
- FIND\_CLUSTERS\_CLEANUP, [9-33](#)
- FIND\_CLUSTERS\_PREP, [9-34](#)
- FIND\_SP, [9-36](#)
- FIND\_SP\_CLEANUP, [9-37](#)
- FIND\_SP\_PREP, [9-38](#)
- GET\_BUILD\_ID, [9-39](#)
- GET\_GEOMETRY\_FROM\_V\_COL, [9-39](#)
- GET\_GEOMETRY\_FROM\_V\_T\_COLS, [9-41](#)
- GET\_LATLONG\_FROM\_V\_COL, [9-42](#), [9-45](#)
- GET\_LATLONG\_FROM\_V\_T\_COLS, [9-43](#)
- GET\_LONG\_LAT\_GEOMETRY, [9-44](#)
- GET\_LONGLAT\_FROM\_V\_T\_COLS, [9-46](#)
- GET\_SCN, [9-47](#)
- GET\_VERSION, [9-47](#)
- GET\_WKTGEOMETRY\_FROM\_V\_COL, [9-48](#)
- GET\_WKTGEOMETRY\_FROM\_V\_T\_COLS, [9-49](#)
- GRANT\_ACCESS, [9-50](#)
- IMP\_EDGE\_TAB\_STATS, [9-51](#)
- IMP\_VERTEX\_TAB\_STATS, [9-52](#)
- PR, [9-54](#)
- PR\_CLEANUP, [9-56](#)
- PR\_PREP, [9-57](#)
- PREPARE\_TEXT\_INDEX, [9-58](#)
- reference information, [9-1](#)

- OPG\_APIS package (*continued*)
  - RENAME\_PG, [9-58](#)
  - SPARSIFY\_GRAPH, [9-59](#)
  - SPARSIFY\_GRAPH\_CLEANUP, [9-61](#)
  - SPARSIFY\_GRAPH\_PREP, [9-62](#)
- OPG\_GRAPHOP package
  - POPULATE\_SKELETON\_TAB, [10-1](#)
  - reference information, [10-1](#)

## P

---

page rank

- calculating, [9-54](#)
- cleanup, [9-56](#)
- preparing to find, [9-57](#)

PGQL (Property Graph Query Language), [6-1](#)

PGX (in-memory Graph server), [4-1](#)

PgXML for Graphs, [8-1](#)

- DeepWalk Algorithm, [8-1](#)
- Pg2vec Algorithm, [8-46](#)
- Supervised GraphWise Algorithm, [8-15](#)
- Unsupervised GraphWise Algorithm, [8-34](#)

POPULATE\_SKELETON\_TAB procedure, [10-1](#)

PR procedure, [9-54](#)

PR\_CLEANUP procedure, [9-56](#)

PR\_PREP procedure, [9-57](#)

PREPARE\_TEXT\_INDEX procedure, [9-58](#)

property graph

- cleanup after sparsifying, [9-61](#)
- clearing (removing data from), [9-10](#)
- cloning, [9-11](#)
- collaborative filtering, [9-4](#), [9-7](#), [9-9](#)
- creating, [9-18](#)
- dropping, [9-26](#)
- dropping view definition, [9-27](#)
- preparing to sparsify, [9-62](#)
- removing text index metadata, [9-11](#)
- renaming, [9-58](#)
- sparsifying, [9-59](#)

property graph access privileges

- granting, [9-50](#)

Property Graph Query Language (PGQL), [6-1](#)

property graph statistics table

- creating, [9-22](#)

property graph support

- getting build ID, [9-39](#)
- getting SCN, [9-47](#)
- getting version, [9-47](#)

## R

---

RENAME\_PG procedure, [9-58](#)

REST Endpoints for Graph Visualization Application, [7-9](#)

- DELETE: Cancel a Query Execution, [7-15](#)



---

**REST Endpoints for Graph Visualization Application (continued)**

- GET: Check a Query Completion, [7-14](#)
- GET: Get User, [7-13](#)
- GET: List Graphs, [7-11](#)
- GET: Retrieve a Query Result, [7-15](#)
- GET: Run a PGQL Query, [7-11](#)
- GET: Run a PGQL Query Asynchronously, [7-14](#)
- POST: Login, [7-10](#)

---

**S**

- shortest path
  - cleanup, [9-37](#)
  - finding, [9-36](#)
  - preparing to find, [9-38](#)
- skeleton table
  - populating, [10-1](#)
- snapshot table
  - creating, [9-19](#)
- SPARSIFY\_GRAPH procedure, [9-59](#)
- SPARSIFY\_GRAPH\_CLEANUP procedure, [9-61](#)
- SPARSIFY\_GRAPH\_PREP procedure, [9-62](#)
- statistics for property graph
  - analyzing, [9-2](#)
- subgraph
  - creating, [9-23](#)

---

**text index**

- on property graph edge table, [9-17](#)
  - on property graph edge table (dropping), [9-26](#)
  - on property graph vertex table, [9-24](#)
  - on property graph vertex table (dropping), [9-27](#)
  - preparing, [9-58](#)
- text index table**
- creating, [9-21](#)
- triangles**
- cleanup after counting, [9-13](#)
  - counting, [9-12](#)
  - counting and renumbering vertices, [9-16](#)
  - estimating the number, [9-28](#)
  - preparing to count, [9-14](#)

---

**V**

- vertex cluster mappings
  - preparing, [9-33](#), [9-34](#)
- vertex table statistics
  - exporting, [9-31](#)
  - importing, [9-52](#)