$\textbf{CX}_O racle$

Release 7.2.1

August 08, 2019

Contents

cx_Oracle is a module that enables access to Oracle Database and conforms to the Python database API specification. This module is currently tested against Oracle Client 19, 18, 12, and 11.2, and Python 2.7, 3.5, 3.6 and 3.7.

cx_Oracle is distributed under an open-source *license* (the BSD license).

Contents:

cx_Oracle 7 Installation

1.1 Overview

To use cx_Oracle 7 with Python and Oracle Database you need:

- Python 2.7 or 3.5 and higher. Older versions of cx_Oracle may work with older versions of Python.

- Oracle client libraries. These can be from the free Oracle Instant Client, or those included in Oracle Database if Python is on the same machine as the database. Oracle client libraries versions 19, 18, 12, and 11.2 are supported on Linux, Windows and macOS. Users have also reported success with other platforms.

- An Oracle Database. Oracle's standard client-server version interoperability allows cx_Oracle to connect to both older and newer databases.

1.2 Quick Start cx_Oracle Installation

- An installation of Python is needed. Python 2.7 and Python 3.5 and higher are supported by cx_Oracle 7.

- Install cx_Oracle from PyPI with:

```
python -m pip install cx_Oracle --upgrade
```

Note: if a binary wheel package is not available for your platform, the source package will be downloaded instead. This will be compiled and the resulting binary installed.

If you are behind a proxy, specify your proxy server:

```
python -m pip install cx_Oracle --proxy=http://proxy.example.com:80 --upgrade
```

- Add Oracle 19, 18, 12 or 11.2 client libraries to your operating system library search path such as `PATH` on Windows or `LD_LIBRARY_PATH` on Linux. On macOS move the files to `~/lib` or `/usr/local/lib`.

    - If your database is on a remote computer, then download and unzip the client libraries from the free Oracle Instant Client "Basic" or "Basic Light" package for your operating system architecture.

      Instant Client on Windows requires an appropriate Microsoft Windows Redistributables. On Linux, the `libaio` (sometimes called `libaio1`) package is needed.

    - Alternatively use the client libraries already available in a locally installed database such as the free Oracle XE release.

Version 19, 18 and 12.2 client libraries can connect to Oracle Database 11.2 or greater. Version 12.1 client libraries can connect to Oracle Database 10.2 or greater. Version 11.2 client libraries can connect to Oracle Database 9.2 or greater.

The database abstraction layer in cx_Oracle is ODPI-C, which means that the ODPI-C installation instructions can be useful to review.

- Create a script like the one below:

```python
# myscript.py

from __future__ import print_function

import cx_Oracle

# Connect as user "hr" with password "welcome" to the "oraclepdb" service running
↪on this computer.
connection = cx_Oracle.connect("hr", "welcome", "localhost/orclpdb")

cursor = connection.cursor()
cursor.execute("""
    SELECT first_name, last_name
    FROM employees
    WHERE department_id = :did AND employee_id > :eid""",
    did = 50,
    eid = 190)
for fname, lname in cursor:
    print("Values:", fname, lname)
```

Locate your Oracle Database username and password, and the database connection string. The connection string is commonly of the format `hostname/servicename`, using the hostname where the database is running, and the service name of the Oracle Database instance.

Substitute your username, password and connection string in the code. Run the Python script, for example:

```
python myscript.py
```

You can learn how to use cx_Oracle from the *API documentation* and samples.

If you run into installation trouble, check out the section on *Troubleshooting*.

1.3 Oracle Client and Oracle Database Interoperability

cx_Oracle requires Oracle Client libraries. The libraries provide the necessary network connectivity to access an Oracle Database instance. They also provide basic and advanced connection management and data features to cx_Oracle.

The simplest way to get Oracle Client libraries is to install the free Oracle Instant Client "Basic" or "Basic Light" package. The libraries are also available in any Oracle Database installation or full Oracle Client installation.

Oracle's standard client-server network interoperability allows connections between different versions of Oracle Client libraries and Oracle Database. For certified configurations see Oracle Support's Doc ID 207303.1. In summary, Oracle Client 19, 18 and 12.2 can connect to Oracle Database 11.2 or greater. Oracle Client 12.1 can connect to Oracle Database 10.2 or greater. Oracle Client 11.2 can connect to Oracle Database 9.2 or greater. The technical restrictions on creating connections may be more flexible. For example Oracle Client 12.2 can successfully connect to Oracle Database 10.2.

cx_Oracle uses the shared library loading mechanism available on each supported platform to load the Oracle Client libraries at runtime. It does not need to be rebuilt for different versions of the libraries. Since a single cx_Oracle binary can use different client versions and also access multiple database versions, it is important your application is tested in your intended release environments. Newer Oracle clients support new features, such as the oraaccess.xml external configuration file available with 12.1 or later clients, session pool improvements, call timeouts with 18 or later clients, and other enhancements.

The cx_Oracle function `clientversion()` can be used to determine which Oracle Client version is in use and the attribute `Connection.version` can be used to determine which Oracle Database version a connection is accessing. These can then be used to adjust application behavior accordingly. Attempts to use some Oracle features that are not supported by a particular client/server combination may result in runtime errors. These include:

- when attempting to access attributes that are not supported by the current Oracle Client library you will get the error "ORA-24315: illegal attribute type"

- when attempting to use implicit results with Oracle Client 11.2 against Oracle Database 12c you will get the error "ORA-29481: Implicit results cannot be returned to client"

- when attempting to get array DML row counts with Oracle Client 11.2 you will get the error "DPI-1050: Oracle Client library must be at version 12.1 or higher"

1.4 Installing cx_Oracle on Linux

This section discusses the generic installation method on Linux. Using Python and cx_Oracle RPM packages on Oracle Linux is discussed in *Installing cx_Oracle RPMs on Oracle Linux*.

1.4.1 Install cx_Oracle

The generic way to install cx_Oracle on Linux is to use Python's Pip package to install cx_Oracle from PyPI:

```
python -m pip install cx_Oracle --upgrade
```

If you are behind a proxy, specify your proxy server:

```
python -m pip install cx_Oracle --proxy=http://proxy.example.com:80 --upgrade
```

This will download and install a pre-compiled binary if one is available for your architecture. If a pre-compiled binary is not available, the source will be downloaded, compiled, and the resulting binary installed. Compiling cx_Oracle requires the *Python.h* header file. If you are using the default python package, this file is in the `python-devel` package or equivalent.

1.4.2 Install Oracle Client

Using cx_Oracle requires Oracle Client libraries to be installed. These provide the necessary network connectivity allowing cx_Oracle to access an Oracle Database instance. Oracle Client versions 19, 18, 12 and 11.2 are supported.

- If your database is on a remote computer, then download the free Oracle Instant Client "Basic" or "Basic Light" package for your operating system architecture. Use the RPM or ZIP packages, based on your preferences.

- Alternatively use the client libraries already available in a locally installed database such as the free Oracle XE release.

Oracle Instant Client Zip Files

To use cx_Oracle with Oracle Instant Client zip files:

1. Download an Oracle 19, 18, 12, or 11.2 "Basic" or "Basic Light" zip file: 64-bit or 32-bit, matching your Python architecture.

2. Unzip the package into a single directory that is accessible to your application. For example:

```
mkdir -p /opt/oracle
cd /opt/oracle
unzip instantclient-basic-linux.x64-19.3.0.0.0dbru.zip
```

3. Install the `libaio` package with sudo or as the root user. For example:

```
sudo yum install libaio
```

On some Linux distributions this package is called `libaio1` instead.

4. If there is no other Oracle software on the machine that will be impacted, permanently add Instant Client to the runtime link path. For example, with sudo or as the root user:

```
sudo sh -c "echo /opt/oracle/instantclient_19_3 > /etc/ld.so.conf.d/oracle-
↪instantclient.conf"
sudo ldconfig
```

Alternatively, set the environment variable `LD_LIBRARY_PATH` to the appropriate directory for the Instant Client version. For example:

```
export LD_LIBRARY_PATH=/opt/oracle/instantclient_19_3:$LD_LIBRARY_PATH
```

5. If you intend to co-locate optional Oracle configuration files such as `tnsnames.ora`, `sqlnet.ora` or `oraaccess.xml` with Instant Client, then put them in the `network/admin` subdirectory. With Instant Client 12.2 or earlier, create this manually. For example:

```
mkdir -p /opt/oracle/instantclient_12_2/network/admin
```

This is the default Oracle configuration directory for executables linked with this Instant Client.

Alternatively, Oracle configuration files can be put in another, accessible directory. Then set the environment variable `TNS_ADMIN` to that directory name.

Oracle Instant Client RPMs

To use cx_Oracle with Oracle Instant Client RPMs:

1. Download an Oracle 19, 18, 12, or 11.2 "Basic" or "Basic Light" RPM: 64-bit or 32-bit, matching your Python architecture.

   Oracle's yum server has Instant Client RPMs for Oracle Linux 7 and Instant Client RPMs for Oracle Linux 6 that can be downloaded without needing a click-through.

2. Install the downloaded RPM with sudo or as the root user. For example:

```
sudo yum install oracle-instantclient19.3-basic-19.3.0.0.0-1.x86_64.rpm
```

Yum will automatically install required dependencies, such as `libaio`.

3. For Instant Client 19, the system library search path is automatically configured during installation.

   For older versions, if there is no other Oracle software on the machine that will be impacted, permanently add Instant Client to the runtime link path. For example, with sudo or as the root user:

```
sudo sh -c "echo /usr/lib/oracle/18.3/client64/lib > /etc/ld.so.conf.d/oracle-
↪instantclient.conf"
sudo ldconfig
```

Alternatively, for version 18 and earlier, every shell running Python will need to have the environment variable `LD_LIBRARY_PATH` set to the appropriate directory for the Instant Client version. For example:

```
export LD_LIBRARY_PATH=/usr/lib/oracle/18.3/client64/lib:$LD_LIBRARY_PATH
```

4. If you intend to co-locate optional Oracle configuration files such as `tnsnames.ora`, `sqlnet.ora` or `oraaccess.xml` with Instant Client, then put them in the `network/admin` subdirectory under `lib/`. With Instant Client 12.2 or earlier, create this manually. For example:

```
sudo mkdir -p /usr/lib/oracle/12.2/client64/lib/network/admin
```

This is the default Oracle configuration directory for executables linked with this Instant Client.

Alternatively, Oracle configuration files can be put in another, accessible directory. Then set the environment variable `TNS_ADMIN` to that directory name.

Local Database or Full Oracle Client

cx_Oracle applications can use Oracle Client 19, 18, 12, or 11.2 libraries from a local Oracle Database or full Oracle Client installation.

The libraries must be either 32-bit or 64-bit, matching your Python architecture.

1. Set required Oracle environment variables by running the Oracle environment script. For example:

```
source /usr/local/bin/oraenv
```

For Oracle Database XE, run:

```
source /u01/app/oracle/product/11.2.0/xe/bin/oracle_env.sh
```

2. Optional Oracle configuration files such as `tnsnames.ora`, `sqlnet.ora` or `oraaccess.xml` can be placed in `$ORACLE_HOME/network/admin`.

Alternatively, Oracle configuration files can be put in another, accessible directory. Then set the environment variable `TNS_ADMIN` to that directory name.

1.5 Installing cx_Oracle RPMs on Oracle Linux

Python and cx_Oracle RPM packages are available from the Oracle Linux yum server. Various versions of Python are easily installed. Using the yum server makes it easy to keep up to date.

Installation instructions are at Oracle Linux for Python Developers.

1.6 Installing cx_Oracle on Windows

1.6.1 Install cx_Oracle

Use Python's Pip package to install cx_Oracle from PyPI:

```
python -m pip install cx_Oracle --upgrade
```

If you are behind a proxy, specify your proxy server:

```
python -m pip install cx_Oracle --proxy=http://proxy.example.com:80 --upgrade
```

This will download and install a pre-compiled binary if one is available for your architecture. If a pre-compiled binary is not available, the source will be downloaded, compiled, and the resulting binary installed.

1.6.2 Install Oracle Client

Using cx_Oracle requires Oracle Client libraries to be installed. These provide the necessary network connectivity allowing cx_Oracle to access an Oracle Database instance. Oracle Client versions 18, 12 and 11.2 are supported.

- If your database is on a remote computer, then download the free Oracle Instant Client "Basic" or "Basic Light" package for your operating system architecture.

- Alternatively use the client libraries already available in a locally installed database such as the free Oracle XE release.

Oracle Instant Client Zip Files

To use cx_Oracle with Oracle Instant Client zip files:

1. Download an Oracle 18, 12, or 11.2 "Basic" or "Basic Light" zip file: 64-bit or 32-bit, matching your Python architecture.

2. Unzip the package into a directory that is accessible to your application. For example unzip `instantclient-basic-windows.x64-18.3.0.0.0dbru.zip` to `C:\oracle\instantclient_18_3`.

3. Add this directory to the `PATH` environment variable. For example, on Windows 7, update `PATH` in Control Panel -> System -> Advanced System Settings -> Advanced -> Environment Variables -> System Variables -> PATH. The Instant Client directory must occur in `PATH` before any other Oracle directories.

   Restart any open command prompt windows.

   To avoid interfering with existing tools that require other Oracle Client versions, instead of updating the system-wide `PATH` variable, you may prefer to write a batch file that sets `PATH`, for example:

```
REM mypy.bat
SET PATH=C:\oracle\instantclient_18_3;%PATH%
python %*
```

   Invoke this batch file every time you want to run python.

   Alternatively use `SET` to change your `PATH` in each command prompt window before you run python.

4. Oracle Instant Client libraries require a Visual Studio redistributable with a 64-bit or 32-bit architecture to match Instant Client's architecture. Each Instant Client version requires a different redistributable version:

   - For Instant Client 18 or 12.2 install VS 2013

   - For Instant Client 12.1 install VS 2010

   - For Instant Client 11.2 install VS 2005 64-bit or VS 2005 32-bit

5. If you intend to co-locate optional Oracle configuration files such as `tnsnames.ora`, `sqlnet.ora` or `oraaccess.xml` with Instant Client, then create a `network\admin` subdirectory, for example `C:\oracle\instantclient_18_3\network\admin`.

   This is the default Oracle configuration directory for executables linked with this Instant Client.

Alternatively, Oracle configuration files can be put in another, accessible directory. Then set the environment variable `TNS_ADMIN` to that directory name.

Local Database or Full Oracle Client

cx_Oracle applications can use Oracle Client 18, 12, or 11.2 libraries libraries from a local Oracle Database or full Oracle Client.

The Oracle libraries must be either 32-bit or 64-bit, matching your Python architecture.

1. Set the environment variable `PATH` to include the path that contains OCI.dll, if it is not already set. For example, on Windows 7, update `PATH` in Control Panel -> System -> Advanced System Settings -> Advanced -> Environment Variables -> System Variables -> PATH.

   Restart any open command prompt windows.

2. Optional Oracle configuration files such as `tnsnames.ora`, `sqlnet.ora` or `oraaccess.xml` can be placed in the `network\admin` subdirectory of the Oracle Database software installation.

   Alternatively, Oracle configuration files can be put in another, accessible directory. Then set the environment variable `TNS_ADMIN` to that directory name.

1.7 Installing cx_Oracle on macOS

1.7.1 Install Python

Make sure you are not using the bundled Python. This has restricted entitlements and will fail to load Oracle client libraries. Instead use Homebrew or Python.org.

1.7.2 Install cx_Oracle

Use Python's Pip package to install cx_Oracle from PyPI:

```
python -m pip install cx_Oracle --upgrade
```

If you are behind a proxy, specify your proxy server:

```
python -m pip install cx_Oracle --proxy=http://proxy.example.com:80 --upgrade
```

The source will be downloaded, compiled, and the resulting binary installed.

1.7.3 Install Oracle Instant Client

cx_Oracle requires Oracle Client libraries, which are found in Oracle Instant Client for macOS. These provide the necessary network connectivity allowing cx_Oracle to access an Oracle Database instance. Oracle Client versions 18, 12 and 11.2 are supported.

To use cx_Oracle with Oracle Instant Client zip files:

1. Download the Oracle 18, 12 or 11.2 "Basic" or "Basic Light" zip file from here. Choose either a 64-bit or 32-bit package, matching your Python architecture.

2. Unzip the package into a single directory that is accessible to your application. For example:

```
mkdir -p /opt/oracle
unzip instantclient-basic-macos.x64-18.1.0.0.0.zip
```

3. Add links to `$HOME/lib` or `/usr/local/lib` to enable applications to find the library. For example:

```
mkdir ~/lib
ln -s /opt/oracle/instantclient_18_1/libclntsh.dylib ~/lib/
```

Alternatively, copy the required OCI libraries. For example:

```
mkdir ~/lib
cp /opt/oracle/instantclient_18_1/{libclntsh.dylib.18.1,libclntshcore.dylib.18.1,
↪libons.dylib,libnnz18.dylib,libociei.dylib} ~/lib/
```

For Instant Client 11.2, the OCI libraries must be copied. For example:

```
mkdir ~/lib
cp /opt/oracle/instantclient_11_2/{libclntsh.dylib.11.1,libnnz11.dylib,libociei.
↪dylib} ~/lib/
```

4. If you intend to co-locate optional Oracle configuration files such as `tnsnames.ora`, `sqlnet.ora` or `oraaccess.xml` with Instant Client, then create a `network/admin` subdirectory, if it does not already exist. For example:

```
mkdir -p /opt/oracle/instantclient_12_2/network/admin
```

This is the default Oracle configuration directory for executables linked with this Instant Client.

Alternatively, Oracle configuration files can be put in another, accessible directory. Then set the environment variable `TNS_ADMIN` to that directory name.

1.8 Installing cx_Oracle without Internet Access

To install cx_Oracle on a computer that is not connected to the internet, download the appropriate cx_Oracle file from PyPI. Transfer this file to the offline computer and install it with:

```
python -m pip install "<file_name>"
```

Then follow the general cx_Oracle platform installation instructions to install Oracle client libraries.

1.9 Install Using GitHub

In order to install using the source on GitHub, use the following commands:

```
git clone https://github.com/oracle/python-cx_Oracle.git cx_Oracle
cd cx_Oracle
git submodule init
git submodule update
python setup.py install
```

Note that if you download a source zip file directly from GitHub then you will also need to download an ODPI-C source zip file and extract it inside the directory called "odpi".

cx_Oracle source code is also available from oss.oracle.com. This can be cloned with:

```
git clone git://oss.oracle.com/git/oracle/python-cx_Oracle.git cx_Oracle
cd cx_Oracle
git submodule init
git submodule update
```

1.10 Install Using Source from PyPI

The source package can be downloaded manually from PyPI and extracted, after which the following commands should be run:

```
python setup.py build
python setup.py install
```

1.11 Upgrading from cx_Oracle 6

Review the release notes for deprecations and modify any affected code.

1.12 Upgrading from cx_Oracle 5

If you are upgrading from cx_Oracle 5 note these installation changes:

- When using Oracle Instant Client, you should not set ORACLE_HOME.

- On Linux, cx_Oracle 6 no longer uses Instant Client RPMs automatically. You must set LD_LIBRARY_PATH or use ldconfig to locate the Oracle client library.

- PyPI no longer allows Windows installers or Linux RPMs to be hosted. Use the supplied cx_Oracle Wheels instead, or use RPMs from Oracle, see *Installing cx_Oracle RPMs on Oracle Linux*.

1.13 Installing cx_Oracle 5.3

If you require cx_Oracle 5.3, download a Windows installer from PyPI or use python -m pip install cx-oracle==5.3 to install from source.

Very old versions of cx_Oracle can be found in the files section at SourceForge.

1.14 Troubleshooting

If installation fails:

- Use option -v with pip. Review your output and logs. Try to install using a different method. **Google anything that looks like an error.** Try some potential solutions.

- Was there a network connection error? Do you need to set the environment variables http_proxy and/or https_proxy? Or try pip install --proxy=http://proxy.example.com:80 cx_Oracle --upgrade?

- If upgrading gave no errors but the old version is still installed, try `pip install cx_Oracle --upgrade --force-reinstall`

- If you do not have access to modify your system version of Python, can you use `pip install cx_Oracle --upgrade --user` or venv?

- Do you get the error "`No module named pip`"? The pip module is builtin to Python from version 2.7.9 but is sometimes removed by the OS. Use the venv module (builtin to Python 3.x) or virtualenv module (Python 2.x) instead.

- Do you get the error "`fatal error: dpi.h: No such file or directory`" when building from source code? Ensure that your source installation has a subdirectory called "odpi" containing files. If missing, review the section on *Install Using GitHub*.

If using cx_Oracle fails:

- Do you get the error "`DPI-1047: Oracle Client library cannot be loaded`"?

   – Check that Python, cx_Oracle and your Oracle Client libraries are all 64-bit or all 32-bit. The `DPI-1047` message will tell you whether the 64-bit or 32-bit Oracle Client is needed for your Python.

   – On Windows, restart your command prompt and use `set PATH` to check the environment variable has the correct Oracle Client listed before any other Oracle directories.

   – On Windows, use the `DIR` command on the directory set in `PATH`. Verify that `OCI.DLL` exists there.

   – On Windows, check that the correct Windows Redistributables have been installed.

   – On Linux, check the `LD_LIBRARY_PATH` environment variable contains the Oracle Client library directory.

   – On macOS, make sure Oracle Instant Client is in `~/lib` or `/usr/local/lib` and that you are not using the bundled Python (use Homebrew or Python.org instead).

- If you have both Python 2 and 3 installed, make sure you are using the correct python and pip (or python3 and pip3) executables.

Module Interface

`cx_Oracle.`**`__future__`**

Special object which contains attributes which control the behavior of cx_Oracle, allowing for opting in for new features. No attributes are currently supported so all attributes will silently ignore being set and will always appear to have the value None.

Note: This method is an extension to the DB API definition.

New in version 6.2.

`cx_Oracle.`**`Binary`**(*string*)

Construct an object holding a binary (long) string value.

`cx_Oracle.`**`clientversion`**()

Return the version of the client library being used as a 5-tuple. The five values are the major version, minor version, update number, patch number and port update number.

Note: This method is an extension to the DB API definition.

`cx_Oracle.`**`Connection`**(*user=None*, *password=None*, *dsn=None*, *mode=cx_Oracle.DEFAULT_AUTH*, *handle=0*, *pool=None*, *threaded=False*, *events=False*, *cclass=None*, *purity=cx_Oracle.ATTR_PURITY_DEFAULT*, *newpassword=None*, *encoding=None*, *nencoding=None*, *edition=None*, *appcontext=[]*, *tag=None*, *matchanytag=False*, *shardingkey=[]*, *supershardingkey=[]*)

`cx_Oracle.`**`connect`**(*user=None*, *password=None*, *dsn=None*, *mode=cx_Oracle.DEFAULT_AUTH*, *handle=0*, *pool=None*, *threaded=False*, *events=False*, *cclass=None*, *purity=cx_Oracle.ATTR_PURITY_DEFAULT*, *newpassword=None*, *encoding=None*, *nencoding=None*, *edition=None*, *appcontext=[]*, *tag=None*, *matchanytag=None*, *shardingkey=[]*, *supershardingkey=[]*)

Constructor for creating a connection to the database. Return a *connection object*. All parameters are optional and can be specified as keyword parameters.

The dsn (data source name) is the TNS entry (from the Oracle names server or tnsnames.ora file) or is a string like the one returned from *makedsn()*. If only one parameter is passed, a connect string is assumed which is to be of the format user/password@dsn, the same format accepted by Oracle applications such as SQL*Plus.

If the mode is specified, it must be one of the *connection authorization modes* which are defined at the module level.

If the handle is specified, it must be of type OCISvcCtx* and is only of use when embedding Python in an application (like PowerBuilder) which has already made the connection.

The pool parameter is expected to be a *session pool object* and the use of this parameter is the equivalent of calling *SessionPool.acquire()*. Parameters not accepted by that method are ignored.

The threaded parameter is expected to be a boolean expression which indicates whether or not Oracle should wrap accesses to connections with a mutex. Doing so in single threaded applications imposes a performance penalty of about 10-15% which is why the default is False.

The events parameter is expected to be a boolean expression which indicates whether or not to initialize Oracle in events mode. This is required for continuous query notification and high availability event notifications.

The cclass parameter is expected to be a string and defines the connection class for database resident connection pooling (DRCP).

The purity parameter is expected to be one of *ATTR_PURITY_NEW*, *ATTR_PURITY_SELF*, or *ATTR_PURITY_DEFAULT*.

The newpassword parameter is expected to be a string if specified and sets the password for the logon during the connection process.

The encoding parameter is expected to be one of the Python standard encodings such as 'UTF-8', if specified, and sets the encoding to use for regular database strings. If not specified, the Oracle environment variable NLS_LANG is used and the Oracle character set is translated to one of the standard encodings, if possible. If the Oracle environment variable NLS_LANG is not set, ASCII is used.

The nencoding parameter is expected to be one of the Python standard encodings such as 'UTF-8', if specified, and sets the encoding to use for national character set database strings. If not specified, the Oracle environment variable NLS_NCHAR is used and the Oracle character set is translated to one of the standard encodings, if possible. If the Oracle environment variable NLS_NCHAR is not used, the Oracle environment variable NLS_LANG is used instead, and if the Oracle environment variable NLS_LANG is not set, ASCII is used.

The edition parameter is expected to be a string if specified and sets the edition to use for the session. It is only relevant if both the client and the database are at least Oracle Database 11.2. If this parameter is used with the cclass parameter the exception "DPI-1058: edition not supported with connection class" will be raised.

The appcontext parameter is expected to be a list of 3-tuples, if specified, and sets the application context for the connection. Application context is available in the database by using the sys_context() PL/SQL method and can be used within a logon trigger as well as any other PL/SQL procedures. Each entry in the list is expected to contain three strings: the namespace, the name and the value.

The tag parameter, if specified, is expected to be a string and will limit the sessions that can be returned from a session pool unless the matchanytag parameter is set to True. In that case sessions with the specified tag will be preferred over others, but if no such sessions are available a session with a different tag may be returned instead. In any case, untagged sessions will always be returned if no sessions with the specified tag are available. Sessions are tagged when they are *released* back to the pool.

The shardingkey and supershardingkey parameters, if specified, are expected to be a sequence of values which will be used to identify the database shard to connect to. Currently only strings are supported for the key values.

cx_Oracle.**Cursor**(*connection*)
    Constructor for creating a cursor. Return a new *cursor object* using the connection.

Note: This method is an extension to the DB API definition.

---

cx_Oracle.**Date**(*year*, *month*, *day*)
> Construct an object holding a date value.

cx_Oracle.**DateFromTicks**(*ticks*)
> Construct an object holding a date value from the given ticks value (number of seconds since the epoch; see the documentation of the standard Python time module for details).

cx_Oracle.**makedsn**(*host*, *port*, *sid=None*, *service_name=None*, *region=None*, *sharding_key=None*, *super_sharding_key=None*)
> Return a string suitable for use as the dsn parameter for `connect()`. This string is identical to the strings that are defined by the Oracle names server or defined in the tnsnames.ora file.

Note: This method is an extension to the DB API definition.

---

cx_Oracle.**SessionPool**(*user=None*, *password=None*, *dsn=None*, *min=1*, *max=2*, *increment=1*, *connectiontype=cx_Oracle.Connection*, *threaded=False*, *getmode=cx_Oracle.SPOOL_ATTRVAL_NOWAIT*, *events=False*, *homogeneous=True*, *externalauth=False*, *encoding=None*, *nencoding=None*, *edition=None*, *timeout=0*, *waitTimeout=0*, *maxLifetimeSession=0*, *sessionCallback=None*)
> Create and return a *session pool object*. Connection pooling in cx_Oracle is handled by Oracle's Session pooling technology. This allows cx_Oracle applications to support features like Application Continuity.
>
> Session pooling creates a pool of available connections to the database, allowing applications to acquire a connection very quickly. It is of primary use in a server where connections are requested in rapid succession and used for a short period of time, for example in a web server.
>
> If the connection type is specified, all calls to `acquire()` will create connection objects of that type, rather than the base type defined at the module level.
>
> The threaded parameter is expected to be a boolean expression which indicates whether Oracle should wrap accesses to connections with a mutex. Doing so in single threaded applications imposes a performance penalty of about 10-15% which is why the default is False.
>
> The events parameter is expected to be a boolean expression which indicates whether or not to initialize Oracle in events mode. This is required for continuous query notification and high availability event notifications.
>
> The homogeneous parameter is expected to be a boolean expression which indicates whether or not to create a homogeneous pool. A homogeneous pool requires that all connections in the pool use the same credentials. As such proxy authentication and external authentication is not possible with a homogeneous pool.
>
> The externalauth parameter is expected to be a boolean expression which indicates whether or not external authentication should be used. External authentication implies that something other than the database is authenticating the user to the database. This includes the use of operating system authentication and Oracle wallets.
>
> The encoding parameter is expected to be one of the Python standard encodings such as 'UTF-8', if specified, and sets the encoding to use for regular database strings. If not specified, the Oracle environment variable NLS_LANG is used and the Oracle character set is translated to one of the standard encodings, if possible. If the Oracle environment variable NLS_LANG is not set, ASCII is used.
>
> The nencoding parameter is expected to be one of the Python standard encodings such as 'UTF-8', if specified, and sets the encoding to use for national character set database strings. If not specified, the Oracle environment variable NLS_NCHAR is used and the Oracle character set is translated to one of the standard encodings, if possible. If the Oracle environment variable NLS_NCHAR is not used, the Oracle environment variable NLS_LANG is used instead, and if the Oracle environment variable NLS_LANG is not set, ASCII is used.

The edition parameter is expected to be a string, if specified, and sets the edition to use for the sessions in the pool. It is only relevant if both the client and the server are at least Oracle Database 11.2.

The timeout parameter is expected to be an integer, if specified, and sets the length of time (in seconds) after which idle sessions in the pool are terminated. Note that termination only occurs when the pool is accessed. The default value of 0 means that no idle sessions are terminated.

The waitTimeout parameter is expected to be an integer, if specified, and sets the length of time (in milliseconds) that the caller should wait for a session to become available in the pool before returning with an error. This value is only used if the getmode parameter is set to the value *cx_Oracle.SPOOL_ATTRVAL_TIMEDWAIT*.

The maxLifetimeSession parameter is expected to be an integer, if specified, and sets the maximum length of time (in seconds) a pooled session may exist. Sessions that are in use will not be closed. They become candidates for termination only when they are released back to the pool and have existed for longer than maxLifetimeSession seconds. Note that termination only occurs when the pool is accessed. The default value is 0 which means that there is no maximum length of time that a pooled session may exist.

The sessionCallback parameter is expected to be either a string or a callable. If the parameter is a string, this refers to a PL/SQL procedure that will be called when *SessionPool.acquire()* requests a tag and that tag does not match the connection's actual tag. Support for the PL/SQL procedure requires Oracle Client libraries 12.2 or later. See the OCI documentation for more information. If the sessionCallback parameter is a callable, however, it will be called when a newly created connection is returned from the pool or when a tag is requested and that tag does not match the connection's actual tag. The callable will be invoked with the connection and the requested tag as its only parameters.

Note: This method is an extension to the DB API definition.

---

cx_Oracle.**Time**(*hour*, *minute*, *second*)
    Construct an object holding a time value.

Note: The time only data type is not supported by Oracle. Calling this function will raise a NotSupportedError exception.

---

cx_Oracle.**TimeFromTicks**(*ticks*)
    Construct an object holding a time value from the given ticks value (number of seconds since the epoch; see the documentation of the standard Python time module for details).

Note: The time only data type is not supported by Oracle. Calling this function will raise a NotSupportedError exception.

---

cx_Oracle.**Timestamp**(*year*, *month*, *day*, *hour*, *minute*, *second*)
    Construct an object holding a time stamp value.

cx_Oracle.**TimestampFromTicks**(*ticks*)
    Construct an object holding a time stamp value from the given ticks value (number of seconds since the epoch; see the documentation of the standard Python time module for details).

2.1 Constants

2.1.1 General

cx_Oracle.**apilevel**
> String constant stating the supported DB API level. Currently '2.0'.

cx_Oracle.**buildtime**
> String constant stating the time when the binary was built.

> **Note:** This constant is an extension to the DB API definition.

---

cx_Oracle.**paramstyle**
> String constant stating the type of parameter marker formatting expected by the interface. Currently 'named' as in 'where name = :name'.

cx_Oracle.**threadsafety**
> Integer constant stating the level of thread safety that the interface supports. Currently 2, which means that threads may share the module and connections, but not cursors. Sharing means that a thread may use a resource without wrapping it using a mutex semaphore to implement resource locking.
>
> Note that in order to make use of multiple threads in a program which intends to connect and disconnect in different threads, the threaded parameter to *connect()* or *SessionPool()* must be true.

cx_Oracle.**version**

cx_Oracle.**__version__**
> String constant stating the version of the module. Currently '7.2.1'.

> **Note:** This attribute is an extension to the DB API definition.

2.1.2 Advanced Queuing: Delivery Modes

These constants are extensions to the DB API definition. They are possible values for the *deliverymode* attribute of the *dequeue options object* passed as the options parameter to the *Connection.deq()* method as well as the *deliverymode* attribute of the *enqueue options object* passed as the options parameter to the *Connection.enq()* method. They are also possible values for the *deliverymode* attribute of the *message properties object* passed as the msgproperties parameter to the *Connection.deq()* and *Connection.enq()* methods.

cx_Oracle.**MSG_BUFFERED**
> This constant is used to specify that enqueue/dequeue operations should enqueue or dequeue buffered messages.

cx_Oracle.**MSG_PERSISTENT**
> This constant is used to specify that enqueue/dequeue operations should enqueue or dequeue persistent messages. This is the default value.

cx_Oracle.**MSG_PERSISTENT_OR_BUFFERED**
> This constant is used to specify that dequeue operations should dequeue either persistent or buffered messages.

2.1.3 Advanced Queuing: Dequeue Modes

These constants are extensions to the DB API definition. They are possible values for the *mode* attribute of the *dequeue options object*. This object is the options parameter for the *Connection.deq()* method.

---

cx_Oracle.**DEQ_BROWSE**
> This constant is used to specify that dequeue should read the message without acquiring any lock on the message (equivalent to a select statement).

cx_Oracle.**DEQ_LOCKED**
> This constant is used to specify that dequeue should read and obtain a write lock on the message for the duration of the transaction (equivalent to a select for update statement).

cx_Oracle.**DEQ_REMOVE**
> This constant is used to specify that dequeue should read the message and update or delete it. This is the default value.

cx_Oracle.**DEQ_REMOVE_NODATA**
> This constant is used to specify that dequeue should confirm receipt of the message but not deliver the actual message content.

2.1.4 Advanced Queuing: Dequeue Navigation Modes

These constants are extensions to the DB API definition. They are possible values for the *navigation* attribute of the *dequeue options object*. This object is the options parameter for the *Connection.deq()* method.

cx_Oracle.**DEQ_FIRST_MSG**
> This constant is used to specify that dequeue should retrieve the first available message that matches the search criteria. This resets the position to the beginning of the queue.

cx_Oracle.**DEQ_NEXT_MSG**
> This constant is used to specify that dequeue should retrieve the next available message that matches the search criteria. If the previous message belongs to a message group, AQ retrieves the next available message that matches the search criteria and belongs to the message group. This is the default.

cx_Oracle.**DEQ_NEXT_TRANSACTION**
> This constant is used to specify that dequeue should skip the remainder of the transaction group and retrieve the first message of the next transaction group. This option can only be used if message grouping is enabled for the current queue.

2.1.5 Advanced Queuing: Dequeue Visibility Modes

These constants are extensions to the DB API definition. They are possible values for the *visibility* attribute of the *dequeue options object*. This object is the options parameter for the *Connection.deq()* method.

cx_Oracle.**DEQ_IMMEDIATE**
> This constant is used to specify that dequeue should perform its work as part of an independent transaction.

cx_Oracle.**DEQ_ON_COMMIT**
> This constant is used to specify that dequeue should be part of the current transaction. This is the default value.

2.1.6 Advanced Queuing: Dequeue Wait Modes

These constants are extensions to the DB API definition. They are possible values for the *wait* attribute of the *dequeue options object*. This object is the options parameter for the *Connection.deq()* method.

cx_Oracle.**DEQ_NO_WAIT**
> This constant is used to specify that dequeue not wait for messages to be available for dequeuing.

cx_Oracle.**DEQ_WAIT_FOREVER**
> This constant is used to specify that dequeue should wait forever for messages to be available for dequeuing. This is the default value.

2.1.7 Advanced Queuing: Enqueue Visibility Modes

These constants are extensions to the DB API definition. They are possible values for the `visibility` attribute of the *enqueue options object*. This object is the options parameter for the `Connection.enq()` method.

cx_Oracle.**ENQ_IMMEDIATE**
> This constant is used to specify that enqueue should perform its work as part of an independent transaction.

cx_Oracle.**ENQ_ON_COMMIT**
> This constant is used to specify that enqueue should be part of the current transaction. This is the default value.

2.1.8 Advanced Queuing: Message States

These constants are extensions to the DB API definition. They are possible values for the `state` attribute of the *message properties object*. This object is the msgproperties parameter for the `Connection.deq()` and `Connection.enq()` methods.

cx_Oracle.**MSG_EXPIRED**
> This constant is used to specify that the message has been moved to the exception queue.

cx_Oracle.**MSG_PROCESSED**
> This constant is used to specify that the message has been processed and has been retained.

cx_Oracle.**MSG_READY**
> This constant is used to specify that the message is ready to be processed.

cx_Oracle.**MSG_WAITING**
> This constant is used to specify that the message delay has not yet been reached.

2.1.9 Advanced Queuing: Other

These constants are extensions to the DB API definition. They are special constants used in advanced queuing.

cx_Oracle.**MSG_NO_DELAY**
> This constant is a possible value for the `delay` attribute of the *message properties object* passed as the msgproperties parameter to the `Connection.deq()` and `Connection.enq()` methods. It specifies that no delay should be imposed and the message should be immediately available for dequeuing. This is also the default value.

cx_Oracle.**MSG_NO_EXPIRATION**
> This constant is a possible value for the `expiration` attribute of the *message properties object* passed as the msgproperties parameter to the `Connection.deq()` and `Connection.enq()` methods. It specifies that the message never expires. This is also the default value.

2.1.10 Connection Authorization Modes

These constants are extensions to the DB API definition. They are possible values for the mode parameter of the `connect()` method.

---

cx_Oracle.**DEFAULT_AUTH**
> This constant is used to specify that default authentication is to take place. This is the default value if no mode is passed at all.
>
> New in version 7.2.

cx_Oracle.**PRELIM_AUTH**
> This constant is used to specify that preliminary authentication is to be used. This is needed for performing database startup and shutdown.

cx_Oracle.**SYSASM**
> This constant is used to specify that SYSASM access is to be acquired.

cx_Oracle.**SYSBKP**
> This constant is used to specify that SYSBACKUP access is to be acquired.

cx_Oracle.**SYSDBA**
> This constant is used to specify that SYSDBA access is to be acquired.

cx_Oracle.**SYSDGD**
> This constant is used to specify that SYSDG access is to be acquired.

cx_Oracle.**SYSKMT**
> This constant is used to specify that SYSKM access is to be acquired.

cx_Oracle.**SYSOPER**
> This constant is used to specify that SYSOPER access is to be acquired.

cx_Oracle.**SYSRAC**
> This constant is used to specify that SYSRAC access is to be acquired.

2.1.11 Database Shutdown Modes

These constants are extensions to the DB API definition. They are possible values for the mode parameter of the *Connection.shutdown()* method.

cx_Oracle.**DBSHUTDOWN_ABORT**
> This constant is used to specify that the caller should not wait for current processing to complete or for users to disconnect from the database. This should only be used in unusual circumstances since database recovery may be necessary upon next startup.

cx_Oracle.**DBSHUTDOWN_FINAL**
> This constant is used to specify that the instance can be truly halted. This should only be done after the database has been shutdown with one of the other modes (except abort) and the database has been closed and dismounted using the appropriate SQL commands.

cx_Oracle.**DBSHUTDOWN_IMMEDIATE**
> This constant is used to specify that all uncommitted transactions should be rolled back and any connected users should be disconnected.

cx_Oracle.**DBSHUTDOWN_TRANSACTIONAL**
> This constant is used to specify that further connections to the database should be prohibited and no new transactions should be allowed. It then waits for all active transactions to complete.

cx_Oracle.**DBSHUTDOWN_TRANSACTIONAL_LOCAL**
> This constant is used to specify that further connections to the database should be prohibited and no new transactions should be allowed. It then waits for only local active transactions to complete.

2.1.12 Event Types

These constants are extensions to the DB API definition. They are possible values for the *Message.type* attribute of the messages that are sent for subscriptions created by the *Connection.subscribe()* method.

cx_Oracle.**EVENT_AQ**
> This constant is used to specify that one or more messages are available for dequeuing on the queue specified when the subscription was created.

cx_Oracle.**EVENT_DEREG**
> This constant is used to specify that the subscription has been deregistered and no further notifications will be sent.

cx_Oracle.**EVENT_NONE**
> This constant is used to specify no information is available about the event.

cx_Oracle.**EVENT_OBJCHANGE**
> This constant is used to specify that a database change has taken place on a table registered with the *Subscription.registerquery()* method.

cx_Oracle.**EVENT_QUERYCHANGE**
> This constant is used to specify that the result set of a query registered with the *Subscription.registerquery()* method has been changed.

cx_Oracle.**EVENT_SHUTDOWN**
> This constant is used to specify that the instance is in the process of being shut down.

cx_Oracle.**EVENT_SHUTDOWN_ANY**
> This constant is used to specify that any instance (when running RAC) is in the process of being shut down.

cx_Oracle.**EVENT_STARTUP**
> This constant is used to specify that the instance is in the process of being started up.

2.1.13 Operation Codes

These constants are extensions to the DB API definition. They are possible values for the operations parameter for the *Connection.subscribe()* method. One or more of these values can be OR'ed together. These values are also used by the *MessageTable.operation* or *MessageQuery.operation* attributes of the messages that are sent.

cx_Oracle.**OPCODE_ALLOPS**
> This constant is used to specify that messages should be sent for all operations.

cx_Oracle.**OPCODE_ALLROWS**
> This constant is used to specify that the table or query has been completely invalidated.

cx_Oracle.**OPCODE_ALTER**
> This constant is used to specify that messages should be sent when a registered table has been altered in some fashion by DDL, or that the message identifies a table that has been altered.

cx_Oracle.**OPCODE_DELETE**
> This constant is used to specify that messages should be sent when data is deleted, or that the message identifies a row that has been deleted.

cx_Oracle.**OPCODE_DROP**
> This constant is used to specify that messages should be sent when a registered table has been dropped, or that the message identifies a table that has been dropped.

cx_Oracle.**OPCODE_INSERT**
> This constant is used to specify that messages should be sent when data is inserted, or that the message identifies a row that has been inserted.

cx_Oracle.**OPCODE_UPDATE**
> This constant is used to specify that messages should be sent when data is updated, or that the message identifies a row that has been updated.

2.1.14 Session Pool Get Modes

These constants are extensions to the DB API definition. They are possible values for the getmode parameter of the *SessionPool()* method.

cx_Oracle.**SPOOL_ATTRVAL_FORCEGET**
> This constant is used to specify that a new connection will be returned if there are no free sessions available in the pool.

cx_Oracle.**SPOOL_ATTRVAL_NOWAIT**
> This constant is used to specify that an exception should be raised if there are no free sessions available in the pool. This is the default value.

cx_Oracle.**SPOOL_ATTRVAL_WAIT**
> This constant is used to specify that the caller should wait until a session is available if there are no free sessions available in the pool.

cx_Oracle.**SPOOL_ATTRVAL_TIMEDWAIT**
> This constant is used to specify that the caller should wait for a period of time (defined by the waitTimeout parameter) for a session to become available before returning with an error.

2.1.15 Session Pool Purity

These constants are extensions to the DB API definition. They are possible values for the purity parameter of the *connect()* method, which is used in database resident connection pooling (DRCP).

cx_Oracle.**ATTR_PURITY_DEFAULT**
> This constant is used to specify that the purity of the session is the default value identified by Oracle (see Oracle's documentation for more information). This is the default value.

cx_Oracle.**ATTR_PURITY_NEW**
> This constant is used to specify that the session acquired from the pool should be new and not have any prior session state.

cx_Oracle.**ATTR_PURITY_SELF**
> This constant is used to specify that the session acquired from the pool need not be new and may have prior session state.

2.1.16 Subscription Grouping Classes

These constants are extensions to the DB API definition. They are possible values for the groupingClass parameter of the *Connection.subscribe()* method.

cx_Oracle.**SUBSCR_GROUPING_CLASS_TIME**
> This constant is used to specify that events are to be grouped by the period of time in which they are received.

2.1.17 Subscription Grouping Types

These constants are extensions to the DB API definition. They are possible values for the groupingType parameter of the `Connection.subscribe()` method.

cx_Oracle.**SUBSCR_GROUPING_TYPE_SUMMARY**
    This constant is used to specify that when events are grouped a summary of the events should be sent instead of the individual events. This is the default value.

cx_Oracle.**SUBSCR_GROUPING_TYPE_LAST**
    This constant is used to specify that when events are grouped the last event that makes up the group should be sent instead of the individual events.

2.1.18 Subscription Namespaces

These constants are extensions to the DB API definition. They are possible values for the namespace parameter of the `Connection.subscribe()` method.

cx_Oracle.**SUBSCR_NAMESPACE_AQ**
    This constant is used to specify that notifications should be sent when a queue has messages available to dequeue.

cx_Oracle.**SUBSCR_NAMESPACE_DBCHANGE**
    This constant is used to specify that database change notification or query change notification messages are to be sent. This is the default value.

2.1.19 Subscription Protocols

These constants are extensions to the DB API definition. They are possible values for the protocol parameter of the `Connection.subscribe()` method.

cx_Oracle.**SUBSCR_PROTO_HTTP**
    This constant is used to specify that notifications will be sent to an HTTP URL when a message is generated. This value is currently not supported.

cx_Oracle.**SUBSCR_PROTO_MAIL**
    This constant is used to specify that notifications will be sent to an e-mail address when a message is generated. This value is currently not supported.

cx_Oracle.**SUBSCR_PROTO_OCI**
    This constant is used to specify that notifications will be sent to the callback routine identified when the subscription was created. It is the default value and the only value currently supported.

cx_Oracle.**SUBSCR_PROTO_SERVER**
    This constant is used to specify that notifications will be sent to a PL/SQL procedure when a message is generated. This value is currently not supported.

2.1.20 Subscription Quality of Service

These constants are extensions to the DB API definition. They are possible values for the qos parameter of the `Connection.subscribe()` method. One or more of these values can be OR'ed together.

cx_Oracle.**SUBSCR_QOS_BEST_EFFORT**
    This constant is used to specify that best effort filtering for query result set changes is acceptable. False positive notifications may be received. This behaviour may be suitable for caching applications.

cx_Oracle.**SUBSCR_QOS_DEREG_NFY**
>   This constant is used to specify that the subscription should be automatically unregistered after the first notification is received.

cx_Oracle.**SUBSCR_QOS_QUERY**
>   This constant is used to specify that notifications should be sent if the result set of the registered query changes. By default no false positive notifications will be generated.

cx_Oracle.**SUBSCR_QOS_RELIABLE**
>   This constant is used to specify that notifications should not be lost in the event of database failure.

cx_Oracle.**SUBSCR_QOS_ROWIDS**
>   This constant is used to specify that the rowids of the inserted, updated or deleted rows should be included in the message objects that are sent.

2.2 Types

cx_Oracle.**BINARY**
>   This type object is used to describe columns in a database that contain binary data. In Oracle this is RAW columns.

cx_Oracle.**BFILE**
>   This type object is used to describe columns in a database that are BFILEs.
>
>   ---
>   **Note:** This type is an extension to the DB API definition.
>
>   ---

cx_Oracle.**BLOB**
>   This type object is used to describe columns in a database that are BLOBs.
>
>   ---
>   **Note:** This type is an extension to the DB API definition.
>
>   ---

cx_Oracle.**BOOLEAN**
>   This type object is used to represent PL/SQL booleans.
>
>   New in version 5.2.1.
>
>   ---
>   **Note:** This type is an extension to the DB API definition. It is only available in Oracle 12.1 and higher and only within PL/SQL. It cannot be used in columns.
>
>   ---

cx_Oracle.**CLOB**
>   This type object is used to describe columns in a database that are CLOBs.
>
>   ---
>   **Note:** This type is an extension to the DB API definition.
>
>   ---

cx_Oracle.**CURSOR**
>   This type object is used to describe columns in a database that are cursors (in PL/SQL these are known as ref cursors).
>
>   ---
>   **Note:** This type is an extension to the DB API definition.
>
>   ---

cx_Oracle.**DATETIME**
> This type object is used to describe columns in a database that are dates.

cx_Oracle.**FIXED_CHAR**
> This type object is used to describe columns in a database that are fixed length strings (in Oracle these is CHAR columns); these behave differently in Oracle than varchar2 so they are differentiated here even though the DB API does not differentiate them.
>
> ---
> **Note:** This attribute is an extension to the DB API definition.
>
> ---

cx_Oracle.**FIXED_NCHAR**
> This type object is used to describe columns in a database that are NCHAR columns in Oracle; these behave differently in Oracle than nvarchar2 so they are differentiated here even though the DB API does not differentiate them.
>
> ---
> **Note:** This type is an extension to the DB API definition.
>
> ---

cx_Oracle.**INTERVAL**
> This type object is used to describe columns in a database that are of type interval day to second.
>
> ---
> **Note:** This type is an extension to the DB API definition.
>
> ---

cx_Oracle.**LOB**
> This type object is the Python type of *BLOB* and *CLOB* data that is returned from cursors.
>
> ---
> **Note:** This type is an extension to the DB API definition.
>
> ---

cx_Oracle.**LONG_BINARY**
> This type object is used to describe columns in a database that are long binary (in Oracle these are LONG RAW columns).
>
> ---
> **Note:** This type is an extension to the DB API definition.
>
> ---

cx_Oracle.**LONG_STRING**
> This type object is used to describe columns in a database that are long strings (in Oracle these are LONG columns).
>
> ---
> **Note:** This type is an extension to the DB API definition.
>
> ---

cx_Oracle.**NATIVE_FLOAT**
> This type object is used to describe columns in a database that are of type binary_double or binary_float.
>
> ---
> **Note:** This type is an extension to the DB API definition.
>
> ---

cx_Oracle.**NATIVE_INT**
> This type object is used to bind integers using Oracle's native integer support, rather than the standard number support.

New in version 5.3.

Note: This type is an extension to the DB API definition.

---

cx_Oracle.**NCHAR**
> This type object is used to describe national character strings (NVARCHAR2) in Oracle.

Note: This type is an extension to the DB API definition.

---

cx_Oracle.**NCLOB**
> This type object is used to describe columns in a database that are NCLOBs.

Note: This type is an extension to the DB API definition.

---

cx_Oracle.**NUMBER**
> This type object is used to describe columns in a database that are numbers.

cx_Oracle.**OBJECT**
> This type object is used to describe columns in a database that are objects.

Note: This type is an extension to the DB API definition.

---

cx_Oracle.**ROWID**
> This type object is used to describe the pseudo column "rowid".

cx_Oracle.**STRING**
> This type object is used to describe columns in a database that are strings (in Oracle this is VARCHAR2 columns).

cx_Oracle.**TIMESTAMP**
> This type object is used to describe columns in a database that are timestamps.

Note: This attribute is an extension to the DB API definition.

2.3 Exceptions

exception cx_Oracle.**Warning**
> Exception raised for important warnings and defined by the DB API but not actually used by cx_Oracle.

exception cx_Oracle.**Error**
> Exception that is the base class of all other exceptions defined by cx_Oracle and is a subclass of the Python StandardError exception (defined in the module exceptions).

exception cx_Oracle.**InterfaceError**
> Exception raised for errors that are related to the database interface rather than the database itself. It is a subclass of Error.

exception cx_Oracle.**DatabaseError**
> Exception raised for errors that are related to the database. It is a subclass of Error.

exception cx_Oracle.**DataError**
> Exception raised for errors that are due to problems with the processed data. It is a subclass of DatabaseError.

exception cx_Oracle.**OperationalError**
> Exception raised for errors that are related to the operation of the database but are not necessarily under the control of the programmer. It is a subclass of DatabaseError.

exception cx_Oracle.**IntegrityError**
> Exception raised when the relational integrity of the database is affected. It is a subclass of DatabaseError.

exception cx_Oracle.**InternalError**
> Exception raised when the database encounters an internal error. It is a subclass of DatabaseError.

exception cx_Oracle.**ProgrammingError**
> Exception raised for programming errors. It is a subclass of DatabaseError.

exception cx_Oracle.**NotSupportedError**
> Exception raised when a method or database API was used which is not supported by the database. It is a subclass of DatabaseError.

2.4 Exception handling

Note: PEP 249 (Python Database API Specification v2.0) says the following about exception values:

> [. . .] The values of these exceptions are not defined. They should give the user a fairly good idea of what went wrong, though. [. . .]

With cx_Oracle every exception object has exactly one argument in the args tuple. This argument is a cx_Oracle._Error object which has the following five read-only attributes.

_Error.**code**
> Integer attribute representing the Oracle error number (ORA-XXXXX).

_Error.**offset**
> Integer attribute representing the error offset when applicable.

_Error.**message**
> String attribute representing the Oracle message of the error. This message is localized by the environment of the Oracle connection.

_Error.**context**
> String attribute representing the context in which the exception was raised.

_Error.**isrecoverable**
> Boolean attribute representing whether the error is recoverable or not. This is False in all cases unless Oracle Database 12.1 is being used on both the server and the client.
>
> New in version 5.3.

This allows you to use the exceptions for example in the following way:

```python
from __future__ import print_function

import cx_Oracle

connection = cx_Oracle.connect("cx_Oracle/dev@localhost/orclpdb")
cursor = connection.cursor()
```

```
try:
    cursor.execute("select 1 / 0 from dual")
except cx_Oracle.DatabaseError as exc:
    error, = exc.args
    print("Oracle-Error-Code:", error.code)
    print("Oracle-Error-Message:", error.message)
```

CHAPTER 3

Connection Object

Note: Any outstanding changes will be rolled back when the connection object is destroyed or closed.

Connection.**__enter__**()
> The entry point for the connection as a context manager. It returns itself.

Note: This method is an extension to the DB API definition.

Connection.**__exit__**()
> The exit point for the connection as a context manager. This will close the connection and roll back any uncommitted transaction.

Note: This method is an extension to the DB API definition.

Connection.**action**
> This write-only attribute sets the action column in the v$session table. It is a string attribute and cannot be set to None – use the empty string instead.

Note: This attribute is an extension to the DB API definition.

Connection.**autocommit**
> This read-write attribute determines whether autocommit mode is on or off. When autocommit mode is on, all statements are committed as soon as they have completed executing.

Note: This attribute is an extension to the DB API definition.

Connection.**begin**([*formatId*, *transactionId*, *branchId*])
> Explicitly begin a new transaction. Without parameters, this explicitly begins a local transaction; otherwise, this

explicitly begins a distributed (global) transaction with the given parameters. See the Oracle documentation for more details.

Note that in order to make use of global (distributed) transactions, the `internal_name` and `external_name` attributes must be set.

Note: This method is an extension to the DB API definition.

Connection.**callTimeout**
This read-write attribute specifies the amount of time (in milliseconds) that a single round-trip to the database may take before a timeout will occur. A value of 0 means that no timeout will take place.

New in version 7.0.

Note: This attribute is an extension to the DB API definition and is only available in Oracle Client 18c and higher.

Connection.**cancel**()
Cancel a long-running transaction.

Note: This method is an extension to the DB API definition.

Connection.**changepassword**(*oldpassword*, *newpassword*)
Change the password of the logon.

Note: This method is an extension to the DB API definition.

Connection.**client_identifier**
This write-only attribute sets the client_identifier column in the v$session table.

Note: This attribute is an extension to the DB API definition.

Connection.**clientinfo**
This write-only attribute sets the client_info column in the v$session table.

Note: This attribute is an extension to the DB API definition.

Connection.**close**()
Close the connection now, rather than whenever __del__ is called. The connection will be unusable from this point forward; an Error exception will be raised if any operation is attempted with the connection.

All open cursors and LOBs created by the connection will be closed and will also no longer be usable.

Internally, references to the connection are held by cursor objects, LOB objects, subscription objects, etc. Once all of these references are released, the connection itself will be closed automatically. Either control references to these related objects carefully or explicitly close connections in order to ensure sufficient resources are available.

Connection.**commit**()
Commit any pending transactions to the database.

Connection.**createlob**(*lobType*)
> Create and return a new temporary *LOB object* of the specified type. The lobType parameter should be one of
> *cx_Oracle.CLOB*, *cx_Oracle.BLOB* or *cx_Oracle.NCLOB*.
>
> New in version 6.2.
>
> ---
> **Note:** This method is an extension to the DB API definition.
> ---

Connection.**current_schema**
> This read-write attribute sets the current schema attribute for the session. Setting this value is the same as exe-
> cuting the SQL statement "ALTER SESSION SET CURRENT_SCHEMA". The attribute is set (and verified)
> on the next call that does a round trip to the server. The value is placed before unqualified database objects in
> SQL statements you then execute.
>
> ---
> **Note:** This attribute is an extension to the DB API definition.
> ---

Connection.**cursor**()
> Return a new *cursor object* using the connection.

Connection.**dbop**
> This write-only attribute sets the database operation that is to be monitored. This can be viewed in the
> DBOP_NAME column of the V$SQL_MONITOR table.
>
> ---
> **Note:** This attribute is an extension to the DB API definition.
> ---

Connection.**deq**(*name*, *options*, *msgproperties*, *payload*)
> Returns a message id after successfully dequeuing a message. The options object can be created using
> *deqoptions()* and the msgproperties object can be created using *msgproperties()*. The payload must
> be an object created using *ObjectType.newobject()*.
>
> New in version 5.3.
>
> Deprecated since version 7.2: Use the methods *Queue.deqOne()* or *Queue.deqMany()* instead.
>
> ---
> **Note:** This method is an extension to the DB API definition.
> ---

Connection.**deqoptions**()
> Returns an object specifying the options to use when dequeuing messages. See *Dequeue Options* for more
> information.
>
> New in version 5.3.
>
> Deprecated since version 7.2: Use the attribute *Queue.deqOptions* instead.
>
> ---
> **Note:** This method is an extension to the DB API definition.
> ---

Connection.**dsn**
> This read-only attribute returns the TNS entry of the database to which a connection has been established.
>
> ---
> **Note:** This attribute is an extension to the DB API definition.
> ---

Connection.**edition**
> This read-only attribute gets the session edition and is only available in Oracle Database 11.2 (both client and server must be at this level or higher for this to work).
>
> New in version 5.3.
>
> ---
>
> **Note:** This attribute is an extension to the DB API definition.
>
> ---

Connection.**encoding**
> This read-only attribute returns the IANA character set name of the character set in use by the Oracle client for regular strings.
>
> ---
>
> **Note:** This attribute is an extension to the DB API definition.
>
> ---

Connection.**enq**(*name*, *options*, *msgproperties*, *payload*)
> Returns a message id after successfully enqueuing a message. The options object can be created using *enqoptions()* and the msgproperties object can be created using *msgproperties()*. The payload must be an object created using *ObjectType.newobject()*.
>
> New in version 5.3.
>
> Deprecated since version 7.2: Use the methods *Queue.enqOne()* or *Queue.enqMany()* instead.
>
> ---
>
> **Note:** This method is an extension to the DB API definition.
>
> ---

Connection.**enqoptions**()
> Returns an object specifying the options to use when enqueuing messages. See *Enqueue Options* for more information.
>
> New in version 5.3.
>
> Deprecated since version 7.2: Use the attribute *Queue.enqOptions* instead.
>
> ---
>
> **Note:** This method is an extension to the DB API definition.
>
> ---

Connection.**external_name**
> This read-write attribute specifies the external name that is used by the connection when logging distributed transactions.
>
> New in version 5.3.
>
> ---
>
> **Note:** This attribute is an extension to the DB API definition.
>
> ---

Connection.**getSodaDatabase**()
> Return a *SodaDatabase* object for Simple Oracle Document Access (SODA). All SODA operations are performed either on the returned SodaDatabase object or from objects created by the returned SodaDatabase object. See here for additional information on SODA.
>
> New in version 7.0.
>
> ---
>
> **Note:** This method is an extension to the DB API definition.
>
> ---

Connection.**gettype**(*name*)
> Return a *type object* given its name. This can then be used to create objects which can be bound to cursors created by this connection.
>
> New in version 5.3.
>
> ---
> **Note:** This method is an extension to the DB API definition.
>
> ---

Connection.**handle**
> This read-only attribute returns the OCI service context handle for the connection. It is primarily provided to facilitate testing the creation of a connection using the OCI service context handle.
>
> ---
> **Note:** This attribute is an extension to the DB API definition.
>
> ---

Connection.**inputtypehandler**
> This read-write attribute specifies a method called for each value that is bound to a statement executed on any cursor associated with this connection. The method signature is handler(cursor, value, arraysize) and the return value is expected to be a variable object or None in which case a default variable object will be created. If this attribute is None, the default behavior will take place for all values bound to statements.
>
> ---
> **Note:** This attribute is an extension to the DB API definition.
>
> ---

Connection.**internal_name**
> This read-write attribute specifies the internal name that is used by the connection when logging distributed transactions.
>
> New in version 5.3.
>
> ---
> **Note:** This attribute is an extension to the DB API definition.
>
> ---

Connection.**ltxid**
> This read-only attribute returns the logical transaction id for the connection. It is used within Oracle Transaction Guard as a means of ensuring that transactions are not duplicated. See the Oracle documentation and the provided sample for more information.
>
> New in version 5.3.

Connection.**maxBytesPerCharacter**
> This read-only attribute returns the maximum number of bytes each character can use for the client character set.
>
> ---
> **Note:** This attribute is an extension to the DB API definition.
>
> ---

Connection.**module**
> This write-only attribute sets the module column in the v$session table. The maximum length for this string is 48 and if you exceed this length you will get ORA-24960.

Connection.**msgproperties**(*payload*, *correlation*, *delay*, *exceptionq*, *expiration*, *priority*)
> Returns an object specifying the properties of messages used in advanced queuing. See *Message Properties* for more information.

Each of the parameters are optional. If specified, they act as a shortcut for setting each of the equivalently named properties.

New in version 5.3.

Changed in version 7.2: Added parameters

Note: This method is an extension to the DB API definition.

Connection.**nencoding**
   This read-only attribute returns the IANA character set name of the national character set in use by the Oracle client.

Note: This attribute is an extension to the DB API definition.

Connection.**outputtypehandler**
   This read-write attribute specifies a method called for each column that is going to be fetched from any cursor associated with this connection. The method signature is handler(cursor, name, defaultType, length, precision, scale) and the return value is expected to be a variable object or None in which case a default variable object will be created. If this attribute is None, the default behavior will take place for all columns fetched from cursors.

Note: This attribute is an extension to the DB API definition.

Connection.**ping**()
   Ping the server which can be used to test if the connection is still active.

Note: This method is an extension to the DB API definition.

Connection.**prepare**()
   Prepare the distributed (global) transaction for commit. Return a boolean indicating if a transaction was actually prepared in order to avoid the error ORA-24756 (transaction does not exist).

Note: This method is an extension to the DB API definition.

Connection.**queue**(*name*, *payloadType=None*)
   Creates a *queue* which is used to enqueue and dequeue messages in Advanced Queueing.

   The name parameter is expected to be a string identifying the queue in which messages are to be enqueued or dequeued.

   The payloadType parameter, if specified, is expected to be an *object type* that identifies the type of payload the queue expects. If not specified, RAW data is enqueued and dequeued.

   New in version 7.2.

Note: This method is an extension to the DB API definition.

Connection.**rollback**()
   Rollback any pending transactions.

Connection.**shutdown**([*mode*])

> Shutdown the database. In order to do this the connection must be connected as *SYSDBA* or *SYSOPER*. Two calls must be made unless the mode specified is *DBSHUTDOWN_ABORT*. An example is shown below:

```python
import cx_Oracle

connection = cx_Oracle.connect(mode = cx_Oracle.SYSDBA)
connection.shutdown(mode = cx_Oracle.DBSHUTDOWN_IMMEDIATE)
cursor = connection.cursor()
cursor.execute("alter database close normal")
cursor.execute("alter database dismount")
connection.shutdown(mode = cx_Oracle.DBSHUTDOWN_FINAL)
```

> **Note:** This method is an extension to the DB API definition.

Connection.**startup**(*force=False*, *restrict=False*)

> Startup the database. This is equivalent to the SQL*Plus command "startup nomount". The connection must be connected as *SYSDBA* or *SYSOPER* with the *PRELIM_AUTH* option specified for this to work. An example is shown below:

```python
import cx_Oracle

connection = cx_Oracle.connect(
        mode=cx_Oracle.SYSDBA | cx_Oracle.PRELIM_AUTH)
connection.startup()
connection = cx_Oracle.connect(mode=cx_Oracle.SYSDBA)
cursor = connection.cursor()
cursor.execute("alter database mount")
cursor.execute("alter database open")
```

> **Note:** This method is an extension to the DB API definition.

Connection.**stmtcachesize**

> This read-write attribute specifies the size of the statement cache. This value can make a significant difference in performance (up to 100x) if you have a small number of statements that you execute repeatedly.

> **Note:** This attribute is an extension to the DB API definition.

Connection.**subscribe**(*namespace=cx_Oracle.SUBSCR_NAMESPACE_DBCHANGE*, *protocol=cx_Oracle.SUBSCR_PROTO_OCI*, *callback=None*, *timeout=0*, *operations=OPCODE_ALLOPS*, *port=0*, *qos=0*, *ipAddress=None*, *groupingClass=0*, *groupingValue=0*, *groupingType=cx_Oracle.SUBSCR_GROUPING_TYPE_SUMMARY*, *name=None*)

> Return a new *subscription object* that receives notifications for events that take place in the database that match the given parameters.

> The namespace parameter specifies the namespace the subscription uses. It can be one of *cx_Oracle.SUBSCR_NAMESPACE_DBCHANGE* or *cx_Oracle.SUBSCR_NAMESPACE_AQ*.

> The protocol parameter specifies the protocol to use when notifications are sent. Currently the only valid value is *cx_Oracle.SUBSCR_PROTO_OCI*.

> The callback is expected to be a callable that accepts a single parameter. A *message object* is passed to this callback whenever a notification is received.

The timeout value specifies that the subscription expires after the given time in seconds. The default value of 0 indicates that the subscription never expires.

The operations parameter enables filtering of the messages that are sent (insert, update, delete). The default value will send notifications for all operations. This parameter is only used when the namespace is set to `cx_Oracle.SUBSCR_NAMESPACE_DBCHANGE`.

The port parameter specifies the listening port for callback notifications from the database server. If not specified, an unused port will be selected by the Oracle Client libraries.

The qos parameter specifies quality of service options. It should be one or more of the following flags, OR'ed together: `cx_Oracle.SUBSCR_QOS_RELIABLE`, `cx_Oracle.SUBSCR_QOS_DEREG_NFY`, `cx_Oracle.SUBSCR_QOS_ROWIDS`, `cx_Oracle.SUBSCR_QOS_QUERY`, `cx_Oracle.SUBSCR_QOS_BEST_EFFORT`.

The ipAddress parameter specifies the IP address (IPv4 or IPv6) in standard string notation to bind for callback notifications from the database server. If not specified, the client IP address will be determined by the Oracle Client libraries.

The groupingClass parameter specifies what type of grouping of notifications should take place. Currently, if set, this value can only be set to the value `cx_Oracle.SUBSCR_GROUPING_CLASS_TIME`, which will group notifications by the number of seconds specified in the groupingValue parameter. The groupingType parameter should be one of the values `cx_Oracle.SUBSCR_GROUPING_TYPE_SUMMARY` (the default) or `cx_Oracle.SUBSCR_GROUPING_TYPE_LAST`.

The name parameter is used to identify the subscription and is specific to the selected namespace. If the namespace parameter is `cx_Oracle.SUBSCR_NAMESPACE_DBCHANGE` then the name is optional and can be any value. If the namespace parameter is `cx_Oracle.SUBSCR_NAMESPACE_AQ`, however, the name must be in the format '<QUEUE_NAME>' for single consumer queues and '<QUEUE_NAME>:<CONSUMER_NAME>' for multiple consumer queues, and identifies the queue that will be monitored for messages. The queue name may include the schema, if needed.

New in version 6.4: The parameters ipAddress, groupingClass, groupingValue, groupingType and name were added.

Note: This method is an extension to the DB API definition.

Note: The subscription can be deregistered in the database by calling the function `unsubscribe()`. If this method is not called and the connection that was used to create the subscription is explicitly closed using the function `close()`, the subscription will not be deregistered in the database.

Connection.**tag**
This read-write attribute initially contains the actual tag of the session that was acquired from a pool by `SessionPool.acquire()`. If the connection was not acquired from a pool or no tagging parameters were specified (tag and matchanytag) when the connection was acquired from the pool, this value will be None. If the value is changed, it must be a string containing name=value pairs like "k1=v1;k2=v2".

If this value is not None when the connection is released back to the pool it will be used to retag the session. This value can be overridden in the call to `SessionPool.release()`.

Note: This attribute is an extension to the DB API definition.

New in version 7.1.

Connection.**tnsentry**
>   This read-only attribute returns the TNS entry of the database to which a connection has been established.

>   ---
>   **Note:** This attribute is an extension to the DB API definition.
>   ---

Connection.**unsubscribe**(*subscr*)
>   Unsubscribe from events in the database that were originally subscribed to using *subscribe()*. The connection used to unsubscribe should be the same one used to create the subscription, or should access the same database and be connected as the same user name.

>   New in version 6.4.

Connection.**username**
>   This read-only attribute returns the name of the user which established the connection to the database.

>   ---
>   **Note:** This attribute is an extension to the DB API definition.
>   ---

Connection.**version**
>   This read-only attribute returns the version of the database to which a connection has been established.

>   ---
>   **Note:** This attribute is an extension to the DB API definition.
>   ---

>   ---
>   **Note:** If you connect to Oracle Database 18 or higher with client libraries 12.2 or lower that you will only receive the base version (such as 18.0.0.0.0) instead of the full version (18.3.0.0.0).
>   ---

Cursor Object

`Cursor.__enter__()`
  The entry point for the cursor as a context manager. It returns itself.

  Note: This method is an extension to the DB API definition.

`Cursor.__exit__()`
  The exit point for the cursor as a context manager. It closes the cursor.

  Note: This method is an extension to the DB API definition.

`Cursor.arraysize`
  This read-write attribute specifies the number of rows to fetch at a time internally and is the default number of rows to fetch with the `fetchmany()` call. It defaults to 100 meaning to fetch 100 rows at a time. Note that this attribute can drastically affect the performance of a query since it directly affects the number of network round trips that need to be performed. This is the reason for setting it to 100 instead of the 1 that the DB API recommends.

`Cursor.bindarraysize`
  This read-write attribute specifies the number of rows to bind at a time and is used when creating variables via `setinputsizes()` or `var()`. It defaults to 1 meaning to bind a single row at a time.

  Note: The DB API definition does not define this attribute.

`Cursor.arrayvar`(*dataType*, *value*[, *size*])
  Create an array variable associated with the cursor of the given type and size and return a *variable object*. The value is either an integer specifying the number of elements to allocate or it is a list and the number of elements allocated is drawn from the size of the list. If the value is a list, the variable is also set with the contents of the list. If the size is not specified and the type is a string or binary, 4000 bytes is allocated. This is needed for passing arrays to PL/SQL (in cases where the list might be empty and the type cannot be determined automatically) or returning arrays from PL/SQL.

Note: The DB API definition does not define this method.

Cursor.**bindnames**()
> Return the list of bind variable names bound to the statement. Note that a statement must have been prepared first.

Note: The DB API definition does not define this method.

Cursor.**bindvars**
> This read-only attribute provides the bind variables used for the last execute. The value will be either a list or a dictionary depending on whether binding was done by position or name. Care should be taken when referencing this attribute. In particular, elements should not be removed or replaced.

Note: The DB API definition does not define this attribute.

Cursor.**callfunc**(*name*, *returnType*, *parameters=[]*, *keywordParameters={}*)
> Call a function with the given name. The return type is specified in the same notation as is required by *setinputsizes()*. The sequence of parameters must contain one entry for each parameter that the function expects. Any keyword parameters will be included after the positional parameters. The result of the call is the return value of the function.

Note: The DB API definition does not define this method.

Note: If you intend to call *Cursor.setinputsizes()* on the cursor prior to making this call, then note that the first item in the parameter list refers to the return value of the function.

Cursor.**callproc**(*name*, *parameters=[]*, *keywordParameters={}*)
> Call a procedure with the given name. The sequence of parameters must contain one entry for each parameter that the procedure expects. The result of the call is a modified copy of the input sequence. Input parameters are left untouched; output and input/output parameters are replaced with possibly new values. Keyword parameters will be included after the positional parameters and are not returned as part of the output sequence.

Note: The DB API definition does not allow for keyword parameters.

Cursor.**close**()
> Close the cursor now, rather than whenever __del__ is called. The cursor will be unusable from this point forward; an Error exception will be raised if any operation is attempted with the cursor.

Cursor.**connection**
> This read-only attribute returns a reference to the connection object on which the cursor was created.

Note: This attribute is an extension to the DB API definition but it is mentioned in PEP 249 as an optional extension.

Cursor.**description**
> This read-only attribute is a sequence of 7-item sequences. Each of these sequences contains information de-

---

scribing one result column: (name, type, display_size, internal_size, precision, scale, null_ok). This attribute will be None for operations that do not return rows or if the cursor has not had an operation invoked via the *execute()* method yet.

The type will be one of the type objects defined at the module level.

Cursor.**execute**(*statement*[, *parameters*], ***keywordParameters*)

Execute a statement against the database. Parameters may be passed as a dictionary or sequence or as keyword parameters. If the parameters are a dictionary, the values will be bound by name and if the parameters are a sequence the values will be bound by position. Note that if the values are bound by position, the order of the variables is from left to right as they are encountered in the statement and SQL statements are processed differently than PL/SQL statements. For this reason, it is generally recommended to bind parameters by name instead of by position.

Parameters passed as a dictionary are name and value pairs. The name maps to the bind variable name used by the statement and the value maps to the Python value you wish bound to that bind variable.

A reference to the statement will be retained by the cursor. If None or the same string object is passed in again, the cursor will execute that statement again without performing a prepare or rebinding and redefining. This is most effective for algorithms where the same statement is used, but different parameters are bound to it (many times). Note that parameters that are not passed in during subsequent executions will retain the value passed in during the last execution that contained them.

For maximum efficiency when reusing an statement, it is best to use the *setinputsizes()* method to specify the parameter types and sizes ahead of time; in particular, None is assumed to be a string of length 1 so any values that are later bound as numbers or dates will raise a TypeError exception.

If the statement is a query, the cursor is returned as a convenience to the caller (so it can be used directly as an iterator over the rows in the cursor); otherwise, None is returned.

Note: The DB API definition does not define the return value of this method.

Cursor.**executemany**(*statement*, *parameters*, *batcherrors=False*, *arraydmlrowcounts=False*)

Prepare a statement for execution against a database and then execute it against all parameter mappings or sequences found in the sequence parameters. The statement is managed in the same way as the *execute()* method manages it. If the size of the buffers allocated for any of the parameters exceeds 2 GB, you will receive the error "DPI-1015: array size of <n> is too large", where <n> varies with the size of each element being allocated in the buffer. If you receive this error, decrease the number of elements in the sequence parameters.

If there are no parameters, or parameters have previously been bound, the number of iterations can be specified as an integer instead of needing to provide a list of empty mappings or sequences.

When true, the batcherrors parameter enables batch error support within Oracle and ensures that the call succeeds even if an exception takes place in one or more of the sequence of parameters. The errors can then be retrieved using *getbatcherrors()*.

When true, the arraydmlrowcounts parameter enables DML row counts to be retrieved from Oracle after the method has completed. The row counts can then be retrieved using *getarraydmlrowcounts()*.

Both the batcherrors parameter and the arraydmlrowcounts parameter can only be true when executing an insert, update, delete or merge statement; in all other cases an error will be raised.

For maximum efficiency, it is best to use the *setinputsizes()* method to specify the parameter types and sizes ahead of time; in particular, None is assumed to be a string of length 1 so any values that are later bound as numbers or dates will raise a TypeError exception.

Cursor.**executemanyprepared**(*numIters*)

Execute the previously prepared and bound statement the given number of times. The variables that are bound

must have already been set to their desired value before this call is made. This method was designed for the case where optimal performance is required as it comes at the expense of compatibility with the DB API.

Note: The DB API definition does not define this method.

Deprecated since version 6.4: Use *executemany()* instead with None for the statement argument and an integer for the parameters argument.

Cursor.**fetchall**()
Fetch all (remaining) rows of a query result, returning them as a list of tuples. An empty list is returned if no more rows are available. Note that the cursor's arraysize attribute can affect the performance of this operation, as internally reads from the database are done in batches corresponding to the arraysize.

An exception is raised if the previous call to *execute()* did not produce any result set or no call was issued yet.

Cursor.**fetchmany**([*numRows=cursor.arraysize*])
Fetch the next set of rows of a query result, returning a list of tuples. An empty list is returned if no more rows are available. Note that the cursor's arraysize attribute can affect the performance of this operation.

The number of rows to fetch is specified by the parameter. If it is not given, the cursor's arraysize attribute determines the number of rows to be fetched. If the number of rows available to be fetched is fewer than the amount requested, fewer rows will be returned.

An exception is raised if the previous call to *execute()* did not produce any result set or no call was issued yet.

Cursor.**fetchone**()
Fetch the next row of a query result set, returning a single tuple or None when no more data is available.

An exception is raised if the previous call to *execute()* did not produce any result set or no call was issued yet.

Cursor.**fetchraw**([*numRows=cursor.arraysize*])
Fetch the next set of rows of a query result into the internal buffers of the defined variables for the cursor. The number of rows actually fetched is returned. This method was designed for the case where optimal performance is required as it comes at the expense of compatibility with the DB API.

An exception is raised if the previous call to *execute()* did not produce any result set or no call was issued yet.

Note: The DB API definition does not define this method.

Cursor.**fetchvars**
This read-only attribute specifies the list of variables created for the last query that was executed on the cursor. Care should be taken when referencing this attribute. In particular, elements should not be removed or replaced.

Note: The DB API definition does not define this attribute.

Cursor.**getarraydmlrowcounts**()
Retrieve the DML row counts after a call to *executemany()* with arraydmlrowcounts enabled. This will return a list of integers corresponding to the number of rows affected by the DML statement for each element of the array passed to *executemany()*.

> **Note:** The DB API definition does not define this method and it is only available for Oracle 12.1 and higher.

Cursor.**getbatcherrors**()
    Retrieve the exceptions that took place after a call to *executemany()* with batcherrors enabled. This will return a list of Error objects, one error for each iteration that failed. The offset can be determined by looking at the offset attribute of the error object.

> **Note:** The DB API definition does not define this method.

Cursor.**getimplicitresults**()
    Return a list of cursors which correspond to implicit results made available from a PL/SQL block or procedure without the use of OUT ref cursor parameters. The PL/SQL block or procedure opens the cursors and marks them for return to the client using the procedure dbms_sql.return_result. Cursors returned in this fashion should not be closed. They will be closed automatically by the parent cursor when it is closed. Closing the parent cursor will invalidate the cursors returned by this method.

    New in version 5.3.

> **Note:** The DB API definition does not define this method and it is only available for Oracle Database 12.1 (both client and server must be at this level or higher). It is most like the DB API method nextset(), but unlike that method (which requires that the next result set overwrite the current result set), this method returns cursors which can be fetched independently of each other.

Cursor.**inputtypehandler**
    This read-write attribute specifies a method called for each value that is bound to a statement executed on the cursor and overrides the attribute with the same name on the connection if specified. The method signature is handler(cursor, value, arraysize) and the return value is expected to be a variable object or None in which case a default variable object will be created. If this attribute is None, the value of the attribute with the same name on the connection is used.

> **Note:** This attribute is an extension to the DB API definition.

Cursor.**__iter__**()
    Returns the cursor itself to be used as an iterator.

> **Note:** This method is an extension to the DB API definition but it is mentioned in PEP 249 as an optional extension.

Cursor.**outputtypehandler**
    This read-write attribute specifies a method called for each column that is to be fetched from this cursor. The method signature is handler(cursor, name, defaultType, length, precision, scale) and the return value is expected to be a variable object or None in which case a default variable object will be created. If this attribute is None, the value of the attribute with the same name on the connection is used instead.

> **Note:** This attribute is an extension to the DB API definition.

Cursor.**parse**(*statement*)
    This can be used to parse a statement without actually executing it (this step is done automatically by Oracle

when a statement is executed).

Note: The DB API definition does not define this method.

Note: You can parse any DML or DDL statement. DDL statements are executed immediately and an implied commit takes place.

Cursor.**prepare**(*statement*[, *tag*])
: This can be used before a call to `execute()` to define the statement that will be executed. When this is done, the prepare phase will not be performed when the call to `execute()` is made with None or the same string object as the statement. If specified the statement will be returned to the statement cache with the given tag. See the Oracle documentation for more information about the statement cache.

Note: The DB API definition does not define this method.

Cursor.**rowcount**
: This read-only attribute specifies the number of rows that have currently been fetched from the cursor (for select statements), that have been affected by the operation (for insert, update, delete and merge statements), or the number of successful executions of the statement (for PL/SQL statements).

Cursor.**rowfactory**
: This read-write attribute specifies a method to call for each row that is retrieved from the database. Ordinarily a tuple is returned for each row but if this attribute is set, the method is called with the tuple that would normally be returned, and the result of the method is returned instead.

Note: The DB API definition does not define this attribute.

Cursor.**scroll**(*value=0*, *mode="relative"*)
: Scroll the cursor in the result set to a new position according to the mode.

  If mode is "relative" (the default value), the value is taken as an offset to the current position in the result set. If set to "absolute", value states an absolute target position. If set to "first", the cursor is positioned at the first row and if set to "last", the cursor is set to the last row in the result set.

  An error is raised if the mode is "relative" or "absolute" and the scroll operation would position the cursor outside of the result set.

  New in version 5.3.

Note: This method is an extension to the DB API definition but it is mentioned in PEP 249 as an optional extension.

Cursor.**scrollable**
: This read-write boolean attribute specifies whether the cursor can be scrolled or not. By default, cursors are not scrollable, as the server resources and response times are greater than nonscrollable cursors. This attribute is checked and the corresponding mode set in Oracle when calling the method `execute()`.

  New in version 5.3.

Note: The DB API definition does not define this attribute.

Cursor.**setinputsizes**(**args*, ***keywordArgs*)
    This can be used before a call to *execute()*, *callfunc()* or *callproc()* to predefine memory areas
    for the operation's parameters. Each parameter should be a type object corresponding to the input that will be
    used or it should be an integer specifying the maximum length of a string parameter. Use keyword parameters
    when binding by name and positional parameters when binding by position. The singleton None can be used as
    a parameter when using positional parameters to indicate that no space should be reserved for that position.

Note: If you plan to use *callfunc()* then be aware that the first parameter in the list refers to the return
value of the function.

Cursor.**setoutputsize**(*size*[, *column*])
    This method does nothing and is retained solely for compatibility with the DB API. The module automatically
    allocates as much space as needed to fetch LONG and LONG RAW columns (or CLOB as string and BLOB as
    bytes).

Cursor.**statement**
    This read-only attribute provides the string object that was previously prepared with *prepare()* or executed
    with *execute()*.

Note: The DB API definition does not define this attribute.

Cursor.**var**(*dataType*[, *size*, *arraysize*, *inconverter*, *outconverter*, *typename*, *encodingErrors*])
    Create a variable with the specified characteristics. This method was designed for use with PL/SQL in/out
    variables where the length or type cannot be determined automatically from the Python object passed in or for
    use in input and output type handlers defined on cursors or connections.

    The dataType parameter specifies the type of data that should be stored in the variable. This should be one of
    the types defined at the module level (such as *cx_Oracle.STRING*) or a Python type that cx_Oracle knows
    how to process (such as str) or an object type returned from the method *Connection.gettype()*.

    The size parameter specifies the length of string and raw variables and is ignored in all other cases. If not
    specified for string and raw variables, the value 4000 is used.

    The arraysize parameter specifies the number of elements the variable will have. If not specified the bind array
    size (usually 1) is used. When a variable is created in an output type handler this parameter should be set to the
    cursor's array size.

    The inconverter and outconverter parameters specify methods used for converting values to/from the database.
    More information can be found in the section on *variable objects*.

    The typename parameter specifies the name of a SQL object type and must be specified when using type
    cx_Oracle.OBJECT unless the type object was passed directly as the first parameter.

    The encodingErrors parameter specifies what should happen when decoding byte strings fetched from the
    database into strings (Python 3) or unicode objects (Python 2). It should be one of the values noted in the
    builtin decode function.

Note: The DB API definition does not define this method.

Variable Objects

Note: The DB API definition does not define this object.

Variable.**actualElements**
> This read-only attribute returns the actual number of elements in the variable. This corresponds to the number of elements in a PL/SQL index-by table for variables that are created using the method *Cursor.arrayvar()*. For all other variables this value will be identical to the attribute *numElements*.

Variable.**bufferSize**
> This read-only attribute returns the size of the buffer allocated for each element in bytes.

Variable.**getvalue**($\lbrack pos=0 \rbrack$)
> Return the value at the given position in the variable. For variables created using the method *Cursor.arrayvar()* the value returned will be a list of each of the values in the PL/SQL index-by table. For variables bound to DML returning statements, the value returned will also be a list corresponding to the returned data for the given execution of the statement (as identified by the pos parameter).

Variable.**inconverter**
> This read-write attribute specifies the method used to convert data from Python to the Oracle database. The method signature is converter(value) and the expected return value is the value to bind to the database. If this attribute is None, the value is bound directly without any conversion.

Variable.**numElements**
> This read-only attribute returns the number of elements allocated in an array, or the number of scalar items that can be fetched into the variable or bound to the variable.

Variable.**outconverter**
> This read-write attribute specifies the method used to convert data from from the Oracle to Python. The method signature is converter(value) and the expected return value is the value to return to Python. If this attribute is None, the value is returned directly without any conversion.

Variable.**setvalue**(*pos*, *value*)
> Set the value at the given position in the variable.

Variable.**size**

> This read-only attribute returns the size of the variable. For strings this value is the size in characters. For all others, this is same value as the attribute bufferSize.

Variable.**type**

> This read-only attribute returns the type of the variable for those variables that bind Oracle objects (it is not present for any other type of variable).

Variable.**values**

> This read-only attribute returns a copy of the value of all actual positions in the variable as a list. This is the equivalent of calling *getvalue()* for each valid position and the length will correspond to the value of the *actualElements* attribute.

CHAPTER 6

SessionPool Object

Note: This object is an extension to the DB API.

Connection pooling in cx_Oracle is handled by SessionPool objects.

SessionPool.**acquire**(*user=None*, *password=None*, *cclass=None*, *purity=cx_Oracle.ATTR_PURITY_DEFAULT*, *tag=None*, *matchanytag=False*, *shardingkey=[]*, *supershardingkey=[]*)
Acquire a connection from the session pool and return a *connection object*.

If the pool is homogeneous, the user and password parameters cannot be specified. If they are, an exception will be raised.

The cclass parameter, if specified, should be a string corresponding to the connection class for database resident connection pooling (DRCP).

The purity parameter is expected to be one of *ATTR_PURITY_NEW*, *ATTR_PURITY_SELF*, or *ATTR_PURITY_DEFAULT*.

The tag parameter, if specified, is expected to be a string with name=value pairs like "k1=v1;k2=v2" and will limit the sessions that can be returned from a session pool unless the matchanytag parameter is set to True. In that case sessions with the specified tag will be preferred over others, but if no such sessions are available a session with a different tag may be returned instead. In any case, untagged sessions will always be returned if no sessions with the specified tag are available. Sessions are tagged when they are *released* back to the pool.

The shardingkey and supershardingkey parameters, if specified, are expected to be a sequence of values which will be used to identify the database shard to connect to. Currently only strings are supported for the key values.

SessionPool.**busy**
This read-only attribute returns the number of sessions currently acquired.

SessionPool.**close**(*force=False*)
Close the session pool now, rather than when the last reference to it is released, which makes it unusable for further work.

If any connections have been acquired and not released back to the pool this method will fail unless the force parameter is set to True.

SessionPool.**drop**(*connection*)

Drop the connection from the pool which is useful if the connection is no longer usable (such as when the session is killed).

SessionPool.**dsn**

This read-only attribute returns the TNS entry of the database to which a connection has been established.

SessionPool.**homogeneous**

This read-write boolean attribute indicates whether the pool is considered homogeneous or not. If the pool is not homogeneous different authentication can be used for each connection acquired from the pool.

SessionPool.**increment**

This read-only attribute returns the number of sessions that will be established when additional sessions need to be created.

SessionPool.**max**

This read-only attribute returns the maximum number of sessions that the session pool can control.

SessionPool.**max_lifetime_session**

This read-write attribute returns the maximum length of time (in seconds) that a pooled session may exist. Sessions that are in use will not be closed. They become candidates for termination only when they are released back to the pool and have existed for longer than max_lifetime_session seconds. Note that termination only occurs when the pool is accessed. A value of 0 means that there is no maximum length of time that a pooled session may exist. This attribute is only available in Oracle Database 12.1.

New in version 5.3.

SessionPool.**min**

This read-only attribute returns the number of sessions with which the session pool was created and the minimum number of sessions that will be controlled by the session pool.

SessionPool.**name**

This read-only attribute returns the name assigned to the session pool by Oracle.

SessionPool.**opened**

This read-only attribute returns the number of sessions currently opened by the session pool.

SessionPool.**release**(*connection*, *tag=None*)

Release the connection back to the pool now, rather than whenever __del__ is called. The connection will be unusable from this point forward; an Error exception will be raised if any operation is attempted with the connection. Any cursors or LOBs created by the connection will also be marked unusable and an Error exception will be raised if any operation is attempted with them.

Internally, references to the connection are held by cursor objects, LOB objects, etc. Once all of these references are released, the connection itself will be released back to the pool automatically. Either control references to these related objects carefully or explicitly release connections back to the pool in order to ensure sufficient resources are available.

If the tag is not None, it is expected to be a string with name=value pairs like "k1=v1;k2=v2" and will override the value in the property *Connection.tag*. If either *Connection.tag* or the tag parameter are not None, the connection will be retagged when it is released back to the pool.

SessionPool.**stmtcachesize**

This read-write attribute specifies the size of the statement cache that will be used as the starting point for any connections that are created by the session pool. Once created, the connection's statement cache size can only be changed by setting the stmtcachesize attribute on the connection itself.

New in version 6.0.

SessionPool.**timeout**

>   This read-write attribute specifies the time (in seconds) after which idle sessions will be terminated in order to maintain an optimum number of open sessions. Note that termination only occurs when the pool is accessed. A value of 0 means that no idle sessions are terminated.

SessionPool.**tnsentry**

>   This read-only attribute returns the TNS entry of the database to which a connection has been established.

SessionPool.**username**

>   This read-only attribute returns the name of the user which established the connection to the database.

SessionPool.**wait_timeout**

>   This read-write attribute specifies the time (in milliseconds) that the caller should wait for a session to become available in the pool before returning with an error. This value is only used if the getmode parameter to *cx_Oracle.SessionPool()* was the value *cx_Oracle.SPOOL_ATTRVAL_TIMEDWAIT*.

>   New in version 6.4.

Subscription Object

Note: This object is an extension the DB API.

Subscription.**callback**
>   This read-only attribute returns the callback that was registered when the subscription was created.

Subscription.**connection**
>   This read-only attribute returns the connection that was used to register the subscription when it was created.

Subscription.**id**
>   This read-only attribute returns the value of `REGID` found in the database view `USER_CHANGE_NOTIFICATION_REGS` or the value of `REG_ID` found in the database view `USER_SUBSCR_REGISTRATIONS`. For AQ subscriptions, the value is 0.

Subscription.**ipAddress**
>   This read-only attribute returns the IP address used for callback notifications from the database server. If not set during construction, this value is None.
>
>   New in version 6.4.

Subscription.**name**
>   This read-only attribute returns the name used to register the subscription when it was created.
>
>   New in version 6.4.

Subscription.**namespace**
>   This read-only attribute returns the namespace used to register the subscription when it was created.

Subscription.**operations**
>   This read-only attribute returns the operations that will send notifications for each table or query that is registered using this subscription.

Subscription.**port**
>   This read-only attribute returns the port used for callback notifications from the database server. If not set during construction, this value is zero.

Subscription.**protocol**
    This read-only attribute returns the protocol used to register the subscription when it was created.

Subscription.**qos**
    This read-only attribute returns the quality of service flags used to register the subscription when it was created.

Subscription.**registerquery**(*statement*[, *args*])
    Register the query for subsequent notification when tables referenced by the query are changed. This behaves similarly to cursor.execute() but only queries are permitted and the args parameter must be a sequence or dictionary. If the qos parameter included the flag cx_Oracle.SUBSCR_QOS_QUERY when the subscription was created, then the ID for the registered query is returned; otherwise, None is returned.

Subscription.**timeout**
    This read-only attribute returns the timeout (in seconds) that was specified when the subscription was created. A value of 0 indicates that there is no timeout.

7.1 Message Objects

Note: This object is created internally when notification is received and passed to the callback procedure specified when a subscription is created.

Message.**consumerName**
    This read-only attribute returns the name of the consumer which generated the notification. It will be populated if the subscription was created with the namespace *cx_Oracle.SUBSCR_NAMESPACE_AQ* and the queue is a multiple consumer queue.

    New in version 6.4.

Message.**dbname**
    This read-only attribute returns the name of the database that generated the notification.

Message.**queries**
    This read-only attribute returns a list of message query objects that give information about query result sets changed for this notification. This attribute will be None if the qos parameter did not include the flag *SUBSCR_QOS_QUERY* when the subscription was created.

Message.**queueName**
    This read-only attribute returns the name of the queue which generated the notification. It will only be populated if the subscription was created with the namespace *cx_Oracle.SUBSCR_NAMESPACE_AQ*.

    New in version 6.4.

Message.**registered**
    This read-only attribute returns whether the subscription which generated this notification is still registered with the database. The subscription is automatically deregistered with the database when the subscription timeout value is reached or when the first notification is sent (when the quality of service flag *cx_Oracle.SUBSCR_QOS_DEREG_NFY* is used).

    New in version 6.4.

Message.**subscription**
    This read-only attribute returns the subscription object for which this notification was generated.

Message.**tables**
    This read-only attribute returns a list of message table objects that give information about the tables changed for this notification. This attribute will be None if the qos parameter included the flag *SUBSCR_QOS_QUERY* when the subscription was created.

Message.**txid**
>   This read-only attribute returns the id of the transaction that generated the notification.

Message.**type**
>   This read-only attribute returns the type of message that has been sent. See the constants section on event types for additional information.

7.2 Message Table Objects

Note: This object is created internally for each table changed when notification is received and is found in the tables attribute of message objects, and the tables attribute of message query objects.

MessageTable.**name**
>   This read-only attribute returns the name of the table that was changed.

MessageTable.**operation**
>   This read-only attribute returns the operation that took place on the table that was changed.

MessageTable.**rows**
>   This read-only attribute returns a list of message row objects that give information about the rows changed on the table. This value is only filled in if the qos parameter to the *Connection.subscribe()* method included the flag *SUBSCR_QOS_ROWIDS*.

7.3 Message Row Objects

Note: This object is created internally for each row changed on a table when notification is received and is found in the rows attribute of message table objects.

MessageRow.**operation**
>   This read-only attribute returns the operation that took place on the row that was changed.

MessageRow.**rowid**
>   This read-only attribute returns the rowid of the row that was changed.

7.4 Message Query Objects

Note: This object is created internally for each query result set changed when notification is received and is found in the queries attribute of message objects.

MessageQuery.**id**
>   This read-only attribute returns the query id of the query for which the result set changed. The value will match the value returned by Subscription.registerquery when the related query was registered.

MessageQuery.**operation**
>   This read-only attribute returns the operation that took place on the query result set that was changed. Valid values for this attribute are *EVENT_DEREG* and *EVENT_QUERYCHANGE*.

`MessageQuery.`**`tables`**
> This read-only attribute returns a list of message table objects that give information about the table changes that caused the query result set to change for this notification.

LOB Objects

Note: This object is an extension the DB API. It is returned whenever Oracle `CLOB`, `BLOB` and `BFILE` columns are fetched.

LOB.**close**()
> Close the LOB. Call this when writing is completed so that the indexes associated with the LOB can be updated – but only if *open()* was called first.

LOB.**fileexists**()
> Return a boolean indicating if the file referenced by the BFILE type LOB exists.

LOB.**getchunksize**()
> Return the chunk size for the internal LOB. Reading and writing to the LOB in chunks of multiples of this size will improve performance.

LOB.**getfilename**()
> Return a two-tuple consisting of the directory alias and file name for a BFILE type LOB.

LOB.**isopen**()
> Return a boolean indicating if the LOB has been opened using the method *open()*.

LOB.**open**()
> Open the LOB for writing. This will improve performance when writing to a LOB in chunks and there are functional or extensible indexes associated with the LOB. If this method is not called, each write will perform an open internally followed by a close after the write has been completed.

LOB.**read**($\left[\textit{offset=1}\left[, \textit{amount}\right]\right]$)
> Return a portion (or all) of the data in the LOB object. Note that the amount and offset are in bytes for BLOB and BFILE type LOBs and in UCS-2 code points for CLOB and NCLOB type LOBs. UCS-2 code points are equivalent to characters for all but supplemental characters. If supplemental characters are in the LOB, the offset and amount will have to be chosen carefully to avoid splitting a character.

LOB.**setfilename**(*dirAlias*, *name*)
> Set the directory alias and name of the BFILE type LOB.

LOB.`size`()

 Returns the size of the data in the LOB object. For BLOB and BFILE type LOBs this is the number of bytes. For CLOB and NCLOB type LOBs this is the number of UCS-2 code points. UCS-2 code points are equivalent to characters for all but supplemental characters.

LOB.`trim`($\big[$*newSize=0*$\big]$)

 Trim the LOB to the new size.

LOB.`write`(*data*$\big[$, *offset=1*$\big]$)

 Write the data to the LOB object at the given offset. The offset is in bytes for BLOB type LOBs and in UCS-2 code points for CLOB and NCLOB type LOBs. UCS-2 code points are equivalent to characters for all but supplemental characters. If supplemental characters are in the LOB, the offset will have to be chosen carefully to avoid splitting a character. Note that if you want to make the LOB value smaller, you must use the *trim()* function.

Object Type Objects

Note: This object is an extension to the DB API. It is returned by the *Connection.gettype()* call and is available as the *Variable.type* for variables containing Oracle objects.

ObjectType ([*sequence*])
: The object type may be called directly and serves as an alternative way of calling *newobject()*.

ObjectType.**attributes**
: This read-only attribute returns a list of the attributes that make up the object type. Each attribute has a name attribute on it.

ObjectType.**iscollection**
: This read-only attribute returns a boolean indicating if the object type refers to a collection or not.

ObjectType.**name**
: This read-only attribute returns the name of the type.

ObjectType.**newobject** ([*sequence*])
: Return a new Oracle object of the given type. This object can then be modified by setting its attributes and then bound to a cursor for interaction with Oracle. If the object type refers to a collection, a sequence may be passed and the collection will be initialized with the items in that sequence.

ObjectType.**schema**
: This read-only attribute returns the name of the schema that owns the type.

9.1 Object Objects

Note: This object is an extension to the DB API. It is returned by the *ObjectType.newobject()* call and can be bound to variables of type *OBJECT*. Attributes can be retrieved and set directly.

Object.**append**(*element*)
> Append an element to the collection object. If no elements exist in the collection, this creates an element at index 0; otherwise, it creates an element immediately following the highest index available in the collection.

Object.**asdict**()
> Return a dictionary where the collection's indexes are the keys and the elements are its values.

Object.**aslist**()
> Return a list of each of the collection's elements in index order.

Object.**copy**()
> Create a copy of the object and return it.

Object.**delete**(*index*)
> Delete the element at the specified index of the collection. If the element does not exist or is otherwise invalid, an error is raised. Note that the indices of the remaining elements in the collection are not changed. In other words, the delete operation creates holes in the collection.

Object.**exists**(*index*)
> Return True or False indicating if an element exists in the collection at the specified index.

Object.**extend**(*sequence*)
> Append all of the elements in the sequence to the collection. This is the equivalent of performing *append()* for each element found in the sequence.

Object.**first**()
> Return the index of the first element in the collection. If the collection is empty, None is returned.

Object.**getelement**(*index*)
> Return the element at the specified index of the collection. If no element exists at that index, an exception is raised.

Object.**last**()
> Return the index of the last element in the collection. If the collection is empty, None is returned.

Object.**next**(*index*)
> Return the index of the next element in the collection following the specified index. If there are no elements in the collection following the specified index, None is returned.

Object.**prev**(*index*)
> Return the index of the element in the collection preceding the specified index. If there are no elements in the collection preceding the specified index, None is returned.

Object.**setelement**(*index*, *value*)
> Set the value in the collection at the specified index to the given value.

Object.**size**()
> Return the number of elements in the collection.

Object.**trim**(*num*)
> Remove the specified number of elements from the end of the collection.

Advanced Queuing

Note: All of these objects are extensions to the DB API.

10.1 Queues

These objects are created using the `Connection.queue()` method and are used to enqueue and dequeue messages.

`Queue.`**`connection`**
> This read-only attribute returns a reference to the connection object on which the queue was created.

`Queue.`**`deqMany`**(*maxMessages*)
> Dequeues up to the specified number of messages from the queue and returns a list of these messages. Each element of the returned list is a *message property* object.

`Queue.`**`deqOne`**()
> Dequeues at most one message from the queue. If a message is dequeued, it will be a *message property* object; otherwise, it will be the value None.

`Queue.`**`deqOptions`**
> This read-only attribute returns a reference to the *options* that will be used when dequeuing messages from the queue.

`Queue.`**`enqOne`**(*message*)
> Enqueues a single message into the queue. The message must be a *message property* object which has had its payload attribute set to a value that the queue supports.

`Queue.`**`enqMany`**(*messages*)
> Enqueues multiple messages into the queue. The messages parameter must be a sequence containing *message property* objects which have all had their payload attribute set to a value that the queue supports.

> Warning: calling this function in parallel on different connections acquired from the same pool may fail due to Oracle bug 29928074. Ensure that this function is not run in parallel, use standalone connections or con-

nections from different pools, or make multiple calls to `Queue.enqOne()` instead. The function `Queue.deqMany()` call is not affected.

Queue.**enqOptions**
> This read-only attribute returns a reference to the *options* that will be used when enqueuing messages into the queue.

Queue.**name**
> This read-only attribute returns the name of the queue.

Queue.**payloadType**
> This read-only attribute returns the object type for payloads that can be enqueued and dequeued. If using a raw queue, this returns the value None.

10.2 Dequeue Options

Note: These objects are used to configure how messages are dequeued from queues. An instance of this object is found in the attribute `Queue.deqOptions`.

DeqOptions.**condition**
> This attribute specifies a boolean expression similar to the where clause of a SQL query. The boolean expression can include conditions on message properties, user data properties and PL/SQL or SQL functions. The default is to have no condition specified.

DeqOptions.**consumername**
> This attribute specifies the name of the consumer. Only messages matching the consumer name will be accessed. If the queue is not set up for multiple consumers this attribute should not be set. The default is to have no consumer name specified.

DeqOptions.**correlation**
> This attribute specifies the correlation identifier of the message to be dequeued. Special pattern-matching characters, such as the percent sign (%) and the underscore (_), can be used. If multiple messages satisfy the pattern, the order of dequeuing is indeterminate. The default is to have no correlation specified.

DeqOptions.**deliverymode**
> This write-only attribute specifies what types of messages should be dequeued. It should be one of the values `MSG_PERSISTENT` (default), `MSG_BUFFERED` or `MSG_PERSISTENT_OR_BUFFERED`.

DeqOptions.**mode**
> This attribute specifies the locking behaviour associated with the dequeue operation. It should be one of the values `DEQ_BROWSE`, `DEQ_LOCKED`, `DEQ_REMOVE` (default), or `DEQ_REMOVE_NODATA`.

DeqOptions.**msgid**
> This attribute specifies the identifier of the message to be dequeued. The default is to have no message identifier specified.

DeqOptions.**navigation**
> This attribute specifies the position of the message that is retrieved. It should be one of the values `DEQ_FIRST_MSG`, `DEQ_NEXT_MSG` (default), or `DEQ_NEXT_TRANSACTION`.

DeqOptions.**transformation**
> This attribute specifies the name of the transformation that must be applied after the message is dequeued from the database but before it is returned to the calling application. The transformation must be created using dbms_transform. The default is to have no transformation specified.

DeqOptions.**visibility**

> This attribute specifies the transactional behavior of the dequeue request. It should be one of the values *DEQ_ON_COMMIT* (default) or *DEQ_IMMEDIATE*. This attribute is ignored when using the *DEQ_BROWSE* mode. Note the value of *autocommit* is always ignored.

DeqOptions.**wait**

> This attribute specifies the time to wait, in seconds, for a message matching the search criteria to become available for dequeuing. One of the values *DEQ_NO_WAIT* or *DEQ_WAIT_FOREVER* can also be used. The default is *DEQ_WAIT_FOREVER*.

10.3 Enqueue Options

Note: These objects are used to configure how messages are enqueued into queues. An instance of this object is found in the attribute *Queue.enqOptions*.

EnqOptions.**deliverymode**

> This write-only attribute specifies what type of messages should be enqueued. It should be one of the values *MSG_PERSISTENT* (default) or *MSG_BUFFERED*.

EnqOptions.**transformation**

> This attribute specifies the name of the transformation that must be applied before the message is enqueued into the database. The transformation must be created using dbms_transform. The default is to have no transformation specified.

EnqOptions.**visibility**

> This attribute specifies the transactional behavior of the enqueue request. It should be one of the values *ENQ_ON_COMMIT* (default) or *ENQ_IMMEDIATE*. Note the value of *autocommit* is ignored.

10.4 Message Properties

Note: These objects are used to identify the properties of messages that are enqueued and dequeued in queues. They are created by the method *Connection.msgproperties()*. They are used by the methods *Queue.enqOne()* and *Queue.enqMany()* and returned by the methods *Queue.deqOne()* and *Queue.deqMany()*.

MessageProperties.**attempts**

> This read-only attribute specifies the number of attempts that have been made to dequeue the message.

MessageProperties.**correlation**

> This attribute specifies the correlation used when the message was enqueued.

MessageProperties.**delay**

> This attribute specifies the number of seconds to delay an enqueued message. Any integer is acceptable but the constant *MSG_NO_DELAY* can also be used indicating that the message is available for immediate dequeuing.

MessageProperties.**deliverymode**

> This read-only attribute specifies the type of message that was dequeued. It will be one of the values *MSG_PERSISTENT* or *MSG_BUFFERED*.

MessageProperties.**enqtime**

> This read-only attribute specifies the time that the message was enqueued.

MessageProperties.**exceptionq**
> This attribute specifies the name of the queue to which the message is moved if it cannot be processed success-fully. Messages are moved if the number of unsuccessful dequeue attempts has exceeded the maximum number of retries or if the message has expired. All messages in the exception queue are in the *MSG_EXPIRED* state. The default value is the name of the exception queue associated with the queue table.

MessageProperties.**expiration**
> This attribute specifies, in seconds, how long the message is available for dequeuing. This attribute is an offset from the delay attribute. Expiration processing requires the queue monitor to be running. Any integer is accepted but the constant *MSG_NO_EXPIRATION* can also be used indicating that the message never expires.

MessageProperties.**msgid**
> This attribute specifies the id of the message in the last queue that generated this message.

MessageProperties.**payload**
> This attribute identifies the payload that will be enqueued or the payload that was dequeued when using a *queue*. When enqueuing, the value is checked to ensure that it conforms to the type expected by that queue. For RAW queues, the value can be a bytes object or a string. If the value is a string it will first be converted to bytes by encoding in the encoding identified by the attribute *Connection.encoding*.

MessageProperties.**priority**
> This attribute specifies the priority of the message. A smaller number indicates a higher priority. The priority can be any integer, including negative numbers. The default value is zero.

MessageProperties.**state**
> This read-only attribute specifies the state of the message at the time of the dequeue. It will be one of the values *MSG_WAITING*, *MSG_READY*, *MSG_PROCESSED* or *MSG_EXPIRED*.

SODA

An overview of Simple Oracle Document Access (SODA) can be found here.

SODA requires Oracle Client 18.3 or higher and Oracle Database 18.1 and higher. The role SODA_APP must be granted to the user.

See this tracking issue for known issues with SODA.

11.1 SODA Database Object

Note: This object is an extension the DB API. It is returned by the method *Connection.getSodaDatabase()*.

SodaDatabase.**createCollection**(*name*, *metadata=None*, *mapMode=False*)
> Creates a SODA collection with the given name and returns a new *SODA collection object*. If you try to create a collection, and a collection with the same name and metadata already exists, then that existing collection is opened without error.
>
> If metadata is specified, it is expected to be a string containing valid JSON or a dictionary that will be transformed into a JSON string. This JSON permits you to specify the configuration of the collection including storage options; specifying the presence or absence of columns for creation timestamp, last modified timestamp and version; whether the collection can store only JSON documents; and methods of key and version generation. The default metadata creates a collection that only supports JSON documents and uses system generated keys. See this collection metadata reference for more information.
>
> If the mapMode parameter is set to True, the new collection is mapped to an existing table instead of creating a table. If a collection is created in this way, dropping the collection will not drop the existing table either.
>
> New in version 7.0.

SodaDatabase.**createDocument**(*content*, *key=None*, *mediaType="application/json"*)
> Creates a *SODA document* usable for SODA write operations. You only need to use this method if your collection requires client-assigned keys or has non-JSON content; otherwise, you can pass your content directly to SODA write operations. SodaDocument attributes 'createdOn', 'lastModified' and 'version' will be None.

The content parameter can be a dictionary or list which will be transformed into a JSON string and then UTF-8 encoded. It can also be a string which will be UTF-8 encoded or it can be a bytes object which will be stored unchanged. If a bytes object is provided and the content is expected to be JSON, note that SODA only supports UTF-8, UTF-16LE and UTF-16BE encodings.

The key parameter should only be supplied if the collection in which the document is to be placed requires client-assigned keys.

The mediaType parameter should only be supplied if the collection in which the document is to be placed supports non-JSON documents and the content for this document is non-JSON. Using a standard MIME type for this value is recommended but any string will be accepted.

New in version 7.0.

SodaDatabase.**getCollectionNames**(*startName=None*, *limit=0*)
Returns a list of the names of collections in the database that match the criteria, in alphabetical order.

If the startName parameter is specified, the list of names returned will start with this value and also contain any names that fall after this value in alphabetical order.

If the limit parameter is specified and is non-zero, the number of collection names returned will be limited to this value.

New in version 7.0.

SodaDatabase.**openCollection**(*name*)
Opens an existing collection with the given name and returns a new *SODA collection object*. If a collection with that name does not exist, None is returned.

New in version 7.0.

11.2 SODA Collection Object

Note: This object is an extension the DB API. It is used to represent SODA collections and is created by methods *SodaDatabase.createCollection()* and *SodaDatabase.openCollection()*.

SodaCollection.**createIndex**(*spec*)
Creates an index on a SODA collection. The spec is expected to be a dictionary or a JSON-encoded string. See this overview for information on indexes in SODA.

Note that a commit should be performed before attempting to create an index.

New in version 7.0.

SodaCollection.**drop**()
Drops the collection from the database, if it exists. Note that if the collection was created with mapMode set to True the underlying table will not be dropped.

A boolean value is returned indicating if the collection was actually dropped.

New in version 7.0.

SodaCollection.**dropIndex**(*name*, *force=False*)
Drops the index with the specified name, if it exists.

The force parameter, if set to True, can be used to force the dropping of an index that the underlying Oracle Database domain index doesn't normally permit. This is only applicable to spatial and JSON search indexes. See here for more information.

A boolean value is returned indicating if the index was actually dropped.

New in version 7.0.

SodaCollection.**find**()
>   This method is used to begin an operation that will act upon documents in the collection. It creates and returns a *SodaOperation object* which is used to specify the criteria and the operation that will be performed on the documents that match that criteria.
>
>   New in version 7.0.

SodaCollection.**getDataGuide**()
>   Returns a *SODA document object* containing property names, data types and lengths inferred from the JSON documents in the collection. It can be useful for exploring the schema of a collection. Note that this method is only supported for JSON-only collections where a JSON search index has been created with the 'dataguide' option enabled. If there are no documents in the collection, None is returned.
>
>   New in version 7.0.

SodaCollection.**insertMany**(*docs*)
>   Inserts a list of documents into the collection at one time. Each of the input documents can be a dictionary or list or an existing *SODA document object*.
>
>   ---
>   **Note:** This method requires Oracle Client 18.5 and higher and is available only as a preview.
>   ---
>
>   New in version 7.2.

SodaCollection.**insertManyAndGet**(*docs*)
>   Similarly to *insertMany()* this method inserts a list of documents into the collection at one time. The only difference is that it returns a list of *SODA Document objects*. Note that for performance reasons the returned documents do not contain the content.
>
>   ---
>   **Note:** This method requires Oracle Client 18.5 and higher.
>   ---
>
>   New in version 7.2.

SodaCollection.**insertOne**(*doc*)
>   Inserts a given document into the collection. The input document can be a dictionary or list or an existing *SODA document object*.
>
>   New in version 7.0.

SodaCollection.**insertOneAndGet**(*doc*)
>   Similarly to *insertOne()* this method inserts a given document into the collection. The only difference is that it returns a *SODA Document object*. Note that for performance reasons the returned document does not contain the content.
>
>   New in version 7.0.

SodaCollection.**metadata**
>   This read-only attribute returns a dictionary containing the metadata that was used to create the collection. See this collection metadata reference for more information.
>
>   New in version 7.0.

SodaCollection.**name**
>   This read-only attribute returns the name of the collection.
>
>   New in version 7.0.

11.3 SODA Document Object

Note: This object is an extension the DB API. It is returned by the methods *SodaDatabase.* *createDocument()*, *SodaOperation.getDocuments()* and *SodaOperation.getOne()* as well as by iterating over *SODA document cursors*.

SodaDoc.**createdOn**
> This read-only attribute returns the creation time of the document in ISO 8601 format. Documents created by *SodaDatabase.createDocument()* or fetched from collections where this attribute is not stored will return None.
>
> New in version 7.0.

SodaDoc.**getContent**()
> Returns the content of the document as a dictionary or list. This method assumes that the content is application/json and will raise an exception if this is not the case. If there is no content, however, None will be returned.
>
> New in version 7.0.

SodaDoc.**getContentAsBytes**()
> Returns the content of the document as a bytes object. If there is no content, however, None will be returned.
>
> New in version 7.0.

SodaDoc.**getContentAsString**()
> Returns the content of the document as a string. If the document encoding is not known, UTF-8 will be used. If there is no content, however, None will be returned.
>
> New in version 7.0.

SodaDoc.**key**
> This read-only attribute returns the unique key assigned to this document. Documents created by *SodaDatabase.createDocument()* may not have a value assigned to them and return None.
>
> New in version 7.0.

SodaDoc.**lastModified**
> This read-only attribute returns the last modified time of the document in ISO 8601 format. Documents created by *SodaDatabase.createDocument()* or fetched from collections where this attribute is not stored will return None.
>
> New in version 7.0.

SodaDoc.**mediaType**
> This read-only attribute returns the media type assigned to the document. By convention this is expected to be a MIME type but no checks are performed on this value. If a value is not specified when calling *SodaDatabase.createDocument()* or the document is fetched from a collection where this component is not stored, the string "application/json" is returned.
>
> New in version 7.0.

SodaDoc.**version**
> This read-only attribute returns the version assigned to this document. Documents created by *SodaDatabase.createDocument()* or fetched from collections where this attribute is not stored will return None.
>
> New in version 7.0.

11.4 SODA Document Cursor Object

Note: This object is an extension the DB API. It is returned by the method `SodaOperation.getCursor()` and implements the iterator protocol. Each iteration will return a *SODA document object*.

SodaDocCursor.**close**()
> Close the cursor now, rather than whenever __del__ is called. The cursor will be unusable from this point forward; an Error exception will be raised if any operation is attempted with the cursor.
>
> New in version 7.0.

11.5 SODA Operation Object

Note: This object is an extension to the DB API. It represents an operation that will be performed on all or some of the documents in a SODA collection. It is created by the method `SodaCollection.find()`.

SodaOperation.**count**()
> Returns a count of the number of documents in the collection that match the criteria. If `skip()` or `limit()` were called on this object, an exception is raised.
>
> New in version 7.0.

SodaOperation.**filter**(*value*)
> Sets a filter specification for complex document queries and ordering of JSON documents. Filter specifications must be provided as a dictionary or JSON-encoded string and can include comparisons, regular expressions, logical and spatial operators, among others. See the overview of SODA filter specifications for more information.
>
> As a convenience, the SodaOperation object is returned so that further criteria can be specified by chaining methods together.
>
> New in version 7.0.

SodaOperation.**getCursor**()
> Returns a *SODA Document Cursor object* that can be used to iterate over the documents that match the criteria.
>
> New in version 7.0.

SodaOperation.**getDocuments**()
> Returns a list of *SODA Document objects* that match the criteria.
>
> New in version 7.0.

SodaOperation.**getOne**()
> Returns a single *SODA Document object* that matches the criteria. Note that if multiple documents match the criteria only the first one is returned.
>
> New in version 7.0.

SodaOperation.**key**(*value*)
> Specifies that the document with the specified key should be returned. This causes any previous calls made to this method and `keys()` to be ignored.
>
> As a convenience, the SodaOperation object is returned so that further criteria can be specified by chaining methods together.

New in version 7.0.

SodaOperation.**keys**(*seq*)
> Specifies that documents that match the keys found in the supplied sequence should be returned. This causes any previous calls made to this method and *key()* to be ignored.

> As a convenience, the SodaOperation object is returned so that further criteria can be specified by chaining methods together.

> New in version 7.0.

SodaOperation.**limit**(*value*)
> Specifies that only the specified number of documents should be returned. This method is only usable for read operations such as *getCursor()* and *getDocuments()*. For write operations, any value set using this method is ignored.

> As a convenience, the SodaOperation object is returned so that further criteria can be specified by chaining methods together.

> New in version 7.0.

SodaOperation.**remove**()
> Removes all of the documents in the collection that match the criteria. The number of documents that have been removed is returned.

> New in version 7.0.

SodaOperation.**replaceOne**(*doc*)
> Replaces a single document in the collection with the specified document. The input document can be a dictionary or list or an existing *SODA document object*. A boolean indicating if a document was replaced or not is returned.

> Currently the method *key()* must be called before this method can be called.

> New in version 7.0.

SodaOperation.**replaceOneAndGet**(*doc*)
> Similarly to *replaceOne()*, this method replaces a single document in the collection with the specified document. The only difference is that it returns a *SODA document object*. Note that for performance reasons the returned document does not contain the content.

> New in version 7.0.

SodaOperation.**skip**(*value*)
> Specifies the number of documents that match the other criteria that will be skipped. This method is only usable for read operations such as *getCursor()* and *getDocuments()*. For write operations, any value set using this method is ignored.

> As a convenience, the SodaOperation object is returned so that further criteria can be specified by chaining methods together.

> New in version 7.0.

SodaOperation.**version**(*value*)
> Specifies that documents with the specified version should be returned. Typically this is used with *key()* to implement optimistic locking, so that the write operation called later does not affect a document that someone else has modified.

> As a convenience, the SodaOperation object is returned so that further criteria can be specified by chaining methods together.

> New in version 7.0.

What's New

12.1 cx_Oracle 6.0

This document contains a summary of the changes in cx_Oracle 6 compared to cx_Oracle 5.3. cx_Oracle 6.0 was released on August 14, 2017. See the *release notes* for complete details.

12.1.1 Highlights

- Has been re-implemented to use the new ODPI-C abstraction layer for Oracle Database. The cx_Oracle API is unchanged. The cx_Oracle design, build and linking process has improved because of ODPI-C.

- Now has Python Wheels available for install. This is made possible by the ODPI-C architecture. Windows installers and Linux RPMs are no longer produced since PyPI no longer supports them.

- Has less code in Python's Global Interpreter Lock, giving better scalability.

- Added support for universal rowids.

- Added support for DML returning of multiple rows.

- Now associates LOB locators to LOB objects so they are not overwritten on database round trips.

12.1.2 Installation Changes

- On Linux, cx_Oracle 6 no longer uses instant client RPMs automatically. You must set LD_LIBRARY_PATH or use ldconfig to locate the Oracle Client library.

- On platforms other than Windows, if ORACLE_HOME is set (in a database or full client installation), remove requirement to set LD_LIBRARY_PATH in order to locate the Oracle Client library (issue 20).

12.1.3 Connection Management Enhancements

- Prevent closing the connection when there are any open statements or LOBs and add new error "DPI-1054: connection cannot be closed when open statements or LOBs exist" when this situation is detected; this is needed to prevent crashes under certain conditions when statements or LOBs are being acted upon while at the same time (in another thread) a connection is being closed; it also prevents leaks of statements and LOBs when a connection is returned to a session pool.

- Added attribute `SessionPool.stmtcachesize` to support getting and setting the default statement cache size for connections in the pool.

- Added attribute `Connection.dbop` to support setting the database operation that is to be monitored.

- Added attribute `Connection.handle` to facilitate testing the creation of a connection using a OCI service context handle.

- Added parameters tag and matchanytag to the `cx_Oracle.connect()` and `SessionPool.acquire()` methods and added parameters tag and retag to the `SessionPool.release()` method in order to support session tagging.

- Added parameter edition to the `cx_Oracle.SessionPool()` method.

- Added parameters region, sharding_key and super_sharding_key to the `cx_Oracle.makedsn()` method to support connecting to a sharded database (new in Oracle Database 12.2).

- Removed requirement that encoding and nencoding both be specified when creating a connection or session pool. The missing value is set to its default value if one of the values is set and the other is not (issue 36).

- Permit use of both string and unicode for Python 2.7 for creating session pools and for changing passwords (issue 23).

12.1.4 Data Type and Data Handling Enhancements

- Added attributes `Variable.actualElements` and `Variable.values` to variables.

- Added support for smallint and float data types in Oracle objects, as requested (issue 4).

- Added support for getting/setting attributes of objects or element values in collections that contain LOBs, BINARY_FLOAT values, BINARY_DOUBLE values and NCHAR and NVARCHAR2 values. The error message for any types that are not supported has been improved as well.

- An exception is no longer raised when a collection is empty for methods `Object.first()` and `Object.last()`. Instead, the value None is returned to be consistent with the methods `Object.next()` and `Object.prev()`.

- Removed requirement for specifying a maximum size when fetching LONG or LONG raw columns. This also allows CLOB, NCLOB, BLOB and BFILE columns to be fetched as strings or bytes without needing to specify a maximum size. The method `Cursor.setoutputsize()` no longer does anything, since ODPI-C automatically manages buffer sizes of LONG and LONG RAW columns.

- Enable temporary LOB caching in order to avoid disk I/O as suggested (issue 10).

12.1.5 Error Handling Enhancements

- Provide improved error message when OCI environment cannot be created, such as when the oraaccess.xml file cannot be processed properly.

- Define exception classes on the connection object in addition to at module scope in order to simplify error handling in multi-connection environments, as specified in the Python DB API.

12.1.6 Test Enhancements

- Reworked test suite and samples so that they are independent of each other and so that the SQL scripts used to create/drop schemas are easily adjusted to use different schema names, if desired.

- Updated DB API test suite stub to support Python 3.

12.1.7 Removals

- Dropped deprecated parameter twophase from the `cx_Oracle.connect()` method. Applications should set the `Connection.internal_name` and `Connection.external_name` attributes instead to a value appropriate to the application.

- Dropped deprecated parameters action, module and clientinfo from the `cx_Oracle.connect()` method. The appcontext parameter should be used instead as shown in this sample.

- Dropped deprecated attribute numbersAsString from *cursor objects*. Use an output type handler instead as shown in this sample.

- Dropped deprecated attributes cqqos and rowids from *subscription objects*. Use the qos attribute instead as shown in this sample.

- Dropped deprecated parameters cqqos and rowids from the `Connection.subscribe()` method. Use the qos parameter instead as shown in this sample.

cx_Oracle Release Notes

13.1 7.x releases

13.1.1 Version 7.2.1 (July 2019)

1) Resolved `MemoryError` exception on Windows when using an output type handler (issue 330).

2) Improved test suite and samples.

3) Improved documentation.

13.1.2 Version 7.2 (July 2019)

1) Updated embedded ODPI-C to version 3.2.

2) Improved AQ support

   - added support for enqueue and dequeue of RAW payloads

   - added support for bulk enqueue and dequeue of messages

   - added new method *Connection.queue()* which creates a new *queue object* in order to simplify AQ usage

   - enhanced method *Connection.msgproperties()* to allow the writable properties of the newly created object to be initialized.

   - the original methods for enqueueing and dequeuing (Connection.deq(), Connection.deqoptions(), Connection.enq() and Connection.enqoptions()) are now deprecated and will be removed in a future version.

3) Removed preview status from existing SODA functionality. See this tracking issue for known issues with SODA.

4) Added support for a preview of SODA bulk insert, available in Oracle Client 18.5 and higher.

5) Added support for setting LOB object attributes, as requested (issue 299).

6) Added mode *cx_Oracle.DEFAULT_AUTH* as requested (issue 293).

7) Added support for setting the LOB prefetch length indicator in order to reduce the number of round trips when fetching LOBs and then subsequently calling `LOB.size()`, `LOB.getchunksize()` or `LOB.read()`.

8) Added support for types BINARY_INTEGER, PLS_INTEGER, ROWID, LONG and LONG RAW when used in PL/SQL.

9) Eliminated deprecation of attribute `Subscription.id`. It is now populated with the value of REGID found in the database view USER_CHANGE_NOTIFICATION_REGS or the value of REG_ID found in the database view USER_SUBSCR_REGISTRATIONS. For AQ subscriptions, the value is 0.

10) Enabled PY_SSIZE_T_CLEAN, as required by Python 3.8 (issue 317).

11) Eliminated memory leak when fetching objects that are atomically null (issue 298).

12) Eliminated bug when processing the string representation of numbers like 1e-08 and 1e-09 (issue 300).

13) Improved error message when the parent cursor is closed before a fetch is attempted from an implicit result cursor.

14) Improved test suite and samples.

15) Improved documentation.

13.1.3 Version 7.1.3 (April 2019)

1) Updated to ODPI-C 3.1.4.

2) Added support for getting the row count for PL/SQL statements (issue 285).

3) Corrected parsing of connect string so that the last @ symbol is searched for instead of the first @ symbol; otherwise, passwords containing an @ symbol will result in the incorrect DSN being extracted (issue 290).

4) Adjusted return value of cursor.callproc() to follow documentation (only positional arguments are returned since the order of keyword parameters cannot be guaranteed in any case) (PR 287).

5) Corrected code getting sample and test parameters by user input when using Python 2.7.

13.1.4 Version 7.1.2 (March 2019)

1) Updated to ODPI-C 3.1.3.

2) Ensured that the strings "-0" and "-0.0" are correctly handled as zero values (issue 274).

3) Eliminated error when startup and shutdown events are generated (ODPI-C issue 102).

4) Enabled the types specified in `Cursor.setinputsizes()` and `Cursor.callfunc()` to be an object type in addition to a Python type, just like in `Cursor.var()`.

5) Reverted changes to return decimal numbers when the numeric precision was too great to be returned accurately as a floating point number. This change had too great an impact on existing functionality and an output type handler can be used to return decimal numbers where that is desirable (issue 279).

6) Eliminated discrepancies in character sets between an external connection handle and the newly created connection handle that references the external connection handle (issue 273).

7) Eliminated memory leak when receiving messages received from subscriptions.

8) Improved test suite and documentation.

13.1.5 Version 7.1.1 (February 2019)

1) Updated to ODPI-C 3.1.2.

2) Corrected code for freeing CQN message objects when multiple queries are registered (ODPI-C issue 96).

3) Improved error messages and installation documentation.

13.1.6 Version 7.1 (February 2019)

1) Updated to ODPI-C 3.1.

2) Improved support for session tagging in session pools by allowing a session callback to be specified when creating a pool via `cx_Oracle.SessionPool()`. Callbacks can be written in Python or in PL/SQL and can be used to improve performance by decreasing round trips to the database needed to set session state. Callbacks written in Python will be invoked for brand new connections (that have never been acquired from the pool before) or when the tag assigned to the connection doesn't match the one that was requested. Callbacks written in PL/SQL will only be invoked when the tag assigned to the connection doesn't match the one that was requested.

3) Added attribute `Connection.tag` to provide access to the actual tag assigned to the connection. Setting this attribute will cause the connection to be retagged when it is released back to the pool.

4) Added support for fetching SYS.XMLTYPE values as strings, as requested (issue 14). Note that this support is limited to the size of VARCHAR2 columns in the database (either 4000 or 32767 bytes).

5) Added support for allowing the typename parameter in method `Cursor.var()` to be None or a valid object type created by the method `Connection.gettype()`, as requested (issue 231).

6) Added support for getting and setting attributes of type RAW on Oracle objects, as requested (ODPI-C issue 72).

7) Added support for performing external authentication with proxy for standalone connections.

8) Added support for mixing integers, floating point and decimal values in data passed to `Cursor.executemany()` (issue 241). The error message raised when a value cannot be converted to an Oracle number was also improved.

9) Adjusted fetching of numeric values so that no precision is lost. If an Oracle number cannot be represented by a Python floating point number a decimal value is automatically returned instead.

10) Corrected handling of multiple calls to method `Cursor.executemany()` where all of the values in one of the columns passed to the first call are all None and a subsequent call has a value other than None in the same column (issue 236).

11) Added additional check for calling `Cursor.setinputsizes()` with an empty dictionary in order to avoid the error "cx_Oracle.ProgrammingError: positional and named binds cannot be intermixed" (issue 199).

12) Corrected handling of values that exceed the maximum value of a plain integer object on Python 2 on Windows (issue 257).

13) Added error message when attempting external authentication with proxy without placing the user name in [] (proxy authentication was previously silently ignored).

14) Exempted additional error messages from forcing a statement to be dropped from the cache (ODPI-C issue 76).

15) Improved dead session detection when using session pools for Oracle Client 12.2 and higher.

16) Ensured that the connection returned from a pool after a failed ping (such as due to a killed session) is not itself marked as needing to be dropped from the pool.

17) Eliminated memory leak under certain circumstances when pooled connections are released back to the pool.

18) Eliminated memory leak when connections are dropped from the pool.

19) Eliminated memory leak when calling `Connection.close()` after fetching collections from the database.

20) Adjusted order in which memory is freed when the last references to SODA collections, documents, document cursors and collection cursors are released, in order to prevent a segfault under certain circumstances.

21) Improved code preventing a statement from binding itself, in order to avoid a potential segfault under certain circumstances.

22) Worked around OCI bug when attempting to free objects that are PL/SQL records, in order to avoid a potential segfault.

23) Improved test suite and samples. Note that default passwords are no longer supplied. New environment variables can be set to specify passwords if desired, or the tests and samples will prompt for the passwords when needed. In addition, a Python script is now available to create and drop the schemas used for the tests and samples.

24) Improved documentation.

13.1.7 Version 7.0 (September 2018)

1) Update to ODPI-C 3.0.

2) Added support for Oracle Client 18 libraries.

3) Added support for SODA (as preview). See *SODA Database*, *SODA Collection* and *SODA Document* for more information.

4) Added support for call timeouts available in Oracle Client 18.1 and higher. See `Connection.callTimeout`.

5) Added support for getting the contents of a SQL collection object as a dictionary, where the keys are the indices of the collection and the values are the elements of the collection. See function `Object.asdict()`.

6) Added support for closing a session pool via the function `SessionPool.close()`. Once closed, further attempts to use any connection that was acquired from the pool will result in the error "DPI-1010: not connected".

7) Added support for setting a LOB attribute of an object with a string or bytes (instead of requiring a temporary LOB to be created).

8) Added support for the packed decimal type used by object attributes with historical types DECIMAL and NU-MERIC (issue 212).

9) On Windows, first attempt to load oci.dll from the same directory as the cx_Oracle module.

10) SQL objects that are created or fetched from the database are now tracked and marked unusable when a connection is closed. This was done in order to avoid a segfault under certain circumstances.

11) Re-enabled dead session detection functionality when using pools for Oracle Client 12.2 and higher in order to handle classes of connection errors such as resource profile limits.

12) Improved error messages when the Oracle Client or Oracle Database need to be at a minimum version in order to support a particular feature.

13) When a connection is used as a context manager, the connection is now closed when the block ends. Attempts to set `cx_Oracle.__future__.ctx_mgr_close` are now ignored.

14) When a DML returning statement is executed, variables bound to it will return an array when calling `Variable.getvalue()`. Attempts to set `cx_Oracle.__future__.dml_ret_array_val` are now ignored.

15) Support for Python 3.4 has been dropped.

16) Added additional test cases.

17) Improved documentation.

13.2 6.x releases

13.2.1 Version 6.4.1 (July 2018)

1) Update to ODPI-C 2.4.2.

   - Avoid buffer overrun due to improper calculation of length byte when converting some negative 39 digit numbers from string to the internal Oracle number format (ODPI-C issue 67).

2) Prevent error "cx_Oracle.ProgrammingError: positional and named binds cannot be intermixed" when calling cursor.setinputsizes() without any parameters and then calling cursor.execute() with named bind parameters (issue 199).

13.2.2 Version 6.4 (July 2018)

1) Update to ODPI-C 2.4.1.

   - Added support for grouping subscriptions. See parameters groupingClass, groupingValue and grouping-Type to function `Connection.subscribe()`.

   - Added support for specifying the IP address a subscription should use instead of having the Oracle Client library determine the IP address on its own. See parameter ipAddress to function `Connection.subscribe()`.

   - Added support for subscribing to notifications when messages are available to dequeue in an AQ queue. The new constant `cx_Oracle.SUBSCR_NAMESPACE_AQ` should be passed to the namespace parameter of function `Connection.subscribe()` in order to get this functionality. Attributes `Message.queueName` and `Message.consumerName` will be populated in notification messages that are received when this namespace is used.

   - Added attribute `Message.registered` to let the notification callback know when the subscription that generated the notification is no longer registered with the database.

   - Added support for timed waits when acquiring a session from a session pool. Use the new constant `cx_Oracle.SPOOL_ATTRVAL_TIMEDWAIT` in the parameter getmode to function `cx_Oracle.SessionPool()` along with the new parameter waitTimeout.

   - Added support for specifying the timeout and maximum lifetime session for session pools when they are created using function `cx_Oracle.SessionPool()`. Previously the pool had to be created before these values could be changed.

   - Avoid memory leak when dequeuing from an empty queue.

   - Ensure that the row count for queries is reset to zero when the statement is executed (issue 193).

   - If the statement should be deleted from the statement cache, first check to see that there is a statement cache currently being used; otherwise, the error "ORA-24300: bad value for mode" will be raised under certain conditions.

2) Added support for using the cursor as a context manager (issue 190).

3) Added parameter "encodingErrors" to function `Cursor.var()` in order to add support for specifying the "errors" parameter to the decode() that takes place internally when fetching strings from the database (issue 162).

4) Added support for specifying an integer for the parameters argument to `Cursor.executemany()`. This allows for batch execution when no parameters are required or when parameters have previously been bound. This replaces Cursor.executemanyprepared() (which is now deprecated and will be removed in cx_Oracle 7).

5) Adjusted the binding of booleans so that outside of PL/SQL they are bound as integers (issue 181).

6) Added support for binding decimal.Decimal values to cx_Oracle.NATIVE_FLOAT as requested (issue 184).

7) Added checks on passing invalid type parameters to methods `Cursor.arrayvar()`, `Cursor.callfunc()` and `Cursor.setinputsizes()`.

8) Corrected handling of cursors and rowids in DML Returning statements.

9) Added sample from David Lapp demonstrating the use of GeoPandas with SDO_GEOMETRY and a sample for demonstrating the use of REF cursors.

10) Adjusted samples and documentation for clarity.

11) Added additional test cases.

13.2.3 Version 6.3.1 (May 2018)

1) Update to ODPI-C 2.3.2.

   - Ensure that a call to unregister a subscription only occurs if the subscription is still registered.

   - Ensure that before a statement is executed any buffers used for DML returning statements are reset.

2) Ensure that behavior with cx_Oracle.__future__.dml_ret_array_val not set or False is the same as the behavior in cx_Oracle 6.2 (issue 176).

13.2.4 Version 6.3 (April 2018)

1) Update to ODPI-C 2.3.1.

   - Fixed binding of LONG data (values exceeding 32KB) when using the function `Cursor.executemany()`.

   - Added code to verify that a CQN subscription is open before permitting it to be used. Error "DPI-1060: subscription was already closed" will now be raised if an attempt is made to use a subscription that was closed earlier.

   - Stopped attempting to unregister a CQN subscription before it was completely registered. This prevents errors encountered during registration from being masked by an error stating that the subscription has not been registered!

   - Added error "DPI-1061: edition is not supported when a new password is specified" to clarify the fact that specifying an edition and a new password at the same time is not supported when creating a connection. Previously the edition value was simply ignored.

   - Improved error message when older OCI client libraries are being used that don't have the method OCI-ClientVersion().

   - Fixed the handling of ANSI types REAL and DOUBLE PRECISION as implemented by Oracle. These types are just subtypes of NUMBER and are different from BINARY_FLOAT and BINARY_DOUBLE (issue 163).

   - Fixed support for true heterogeneous session pools that use different user/password combinations for each session acquired from the pool.

  • Added error message indicating that setting either of the parameters arraydmlrowcounts and batcherrors to
    True in `Cursor.executemany()` is only supported with insert, update, delete and merge statements.

2) Fixed support for getting the OUT values of bind variables bound to a DML Returning statement when call-
   ing the function `Cursor.executemany()`. Note that the attribute dml_ret_array_val in `cx_Oracle.`
   `__future__` must be set to True first.

3) Added support for binding integers and floats as cx_Oracle.NATIVE_FLOAT.

4) A `cx_Oracle._Error` object is now the value of all cx_Oracle exceptions raised by cx_Oracle. (issue 51).

5) Added support for building cx_Oracle with a pre-compiled version of ODPI-C, as requested (issue 103).

6) Default values are now provided for all parameters to `cx_Oracle.SessionPool()`.

7) Improved error message when an unsupported Oracle type is encountered.

8) The Python GIL is now prevented from being held while performing a round trip for the call to get the attribute
   `Connection.version` (issue 158).

9) Added check for the validity of the year for Python 2.x since it doesn't do that itself like Python 3.x does (issue
   166).

10) Adjusted documentation to provide additional information on the use of `Cursor.executemany()` as re-
    quested (issue 153).

11) Adjusted documentation to state that batch errors and array DML row counts can only be used with insert,
    update, delete and merge statements (issue 31).

12) Updated tutorial to import common connection information from files in order to make setup a bit more generic.

13.2.5 Version 6.2.1 (March 2018)

1) Make sure cxoModule.h is included in the source archive (issue 155).

13.2.6 Version 6.2 (March 2018)

1) Update to ODPI-C 2.2.1.

   • eliminate error "DPI-1054: connection cannot be closed when open statements or LOBs exist" (issue 138).

   • avoid a round trip to the database when a connection is released back to the pool by preventing a rollback
     from being called when no transaction is in progress.

   • improve error message when the use of bind variables is attempted with DLL statements, which is not
     supported by Oracle.

   • if an Oracle object is retrieved from an attribute of another Oracle object or a collection, prevent the
     "owner" from being destroyed until the object that was retrieved has itself been destroyed.

   • correct handling of boundary numbers 1e126 and -1e126

   • eliminate memory leak when calling `Connection.enq()` and `Connection.deq()`

   • eliminate memory leak when setting NCHAR and NVARCHAR attributes of objects.

   • eliminate memory leak when fetching collection objects from the database.

2) Added support for creating a temporary CLOB, BLOB or NCLOB via the method `Connection.`
   `createlob()`.

3) Added support for binding a LOB value directly to a cursor.

4) Added support for closing the connection when reaching the end of a `with` code block controlled by the connection as a context manager, but in a backwards compatible way (issue 113). See `cx_Oracle.__future__` for more information.

5) Reorganized code to simplify continued maintenance and consolidate transformations to/from Python objects.

6) Ensure that the number of elements in the array is not lost when the buffer size is increased to accommodate larger strings.

7) Corrected support in Python 3.x for cursor.parse() by permitting a string to be passed, instead of incorrectly requiring a bytes object.

8) Eliminate reference leak with LOBs acquired from attributes of objects or elements of collections.

9) Eliminate reference leak when extending an Oracle collection.

10) Documentation improvements.

11) Added test cases to the test suite.

13.2.7 Version 6.1 (December 2017)

1) Update to ODPI-C 2.1.

   - Support was added for accessing sharded databases via sharding keys (new in Oracle 12.2). NOTE: the underlying OCI library has a bug when using standalone connections. There is a small memory leak proportional to the number of connections created/dropped. There is no memory leak when using session pools, which is recommended.

   - Added options for authentication with SYSBACKUP, SYSDG, SYSKM and SYSRAC, as requested (issue 101).

   - Attempts to release statements or free LOBs after the connection has been closed (by, for example, killing the session) are now prevented.

   - An error message was added when specifying an edition and a connection class since this combination is not supported.

   - Attempts to close the session for connections created with an external handle are now prevented.

   - Attempting to ping a database earlier than 10g results in ORA-1010: invalid OCI operation, but that implies a response from the database and therefore a successful ping, so treat it that way! (see https://github.com/rana/ora/issues/224 for more information).

   - Support was added for converting numeric values in an object type attribute to integer and text, as requested (ODPI-C issue 35).

   - Setting attributes `DeqOptions.msgId` and `MessageProperties.msgId` now works as expected.

   - The overflow check when using double values (Python floats) as input to float attributes of objects or elements of collections was removed as it didn't work anyway and is a well-known issue that cannot be prevented without removing desired functionality. The developer should ensure that the source value falls within the limits of floats, understand the consequent precision loss or use a different data type.

   - Variables of string/raw types are restricted to 2 bytes less than 1 GB (1,073,741,822 bytes), since OCI cannot handle more than that currently.

   - Support was added for identifying the id of the transaction which spawned a CQN subscription message, as requested (ODPI-C issue 32).

   - Corrected use of subscription port number (issue 115).

   - Problems reported with the usage of FormatMessage() on Windows were addressed (ODPI-C issue 47).

- On Windows, if oci.dll cannot be loaded because it is the wrong architecture (32-bit vs 64-bit), attempt to find the offending DLL and include the full path of the DLL in the message, as suggested. (ODPI-C issue 49).

- Force OCI prefetch to always use the value 2; the OCI default is 1 but setting the ODPI-C default to 2 ensures that single row fetches don't require an extra round trip to determine if there are more rows to fetch; this change also reduces the potential memory consumption when fetchArraySize was set to a large value and also avoids performance issues discovered with larger values of prefetch.

2) Fix build with PyPy 5.9.0-alpha0 in libpython mode (PR 54).

3) Ensure that the edition is passed through to the database when a session pool is created.

4) Corrected handling of Python object references when an invalid keyword parameter is passed to *cx_Oracle.SessionPool()*.

5) Corrected handling of *Connection.handle* and the handle parameter to *cx_Oracle.connect()* on Windows.

6) Documentation improvements.

7) Added test cases to the test suite.

13.2.8 Version 6.0.3 (November 2017)

1) Update to ODPI-C 2.0.3.

- Prevent use of uninitialized data in certain cases (issue 77).

- Attempting to ping a database earlier than 10g results in error "ORA-1010: invalid OCI operation", but that implies a response from the database and therefore a successful ping, so treat it that way!

- Correct handling of conversion of some numbers to NATIVE_FLOAT.

- Prevent use of NaN with Oracle numbers since it produces corrupt data (issue 91).

- Verify that Oracle objects bound to cursors, fetched from cursors, set in object attributes or appended to collection objects are of the correct type.

- Correct handling of NVARCHAR2 when used as attributes of Oracle objects or as elements of collections.

2) Ensure that a call to setinputsizes() with an invalid type prior to a call to executemany() does not result in a type error, but instead gracefully ignores the call to setinputsizes() as required by the DB API (issue 75).

3) Check variable array size when setting variable values and raise IndexError, as is already done for getting variable values.

13.2.9 Version 6.0.2 (August 2017)

1) Update to ODPI-C 2.0.2.

- Don't prevent connection from being explicitly closed when a fatal error has taken place (issue 67).

- Correct handling of objects when dynamic binding is performed.

- Process deregistration events without an error.

- Eliminate memory leak when creating objects.

2) Added missing type check to prevent coercion of decimal to float (issue 68).

3) On Windows, sizeof(long) = 4, not 8, which meant that integers between 10 and 18 digits were not converted to Python correctly (issue 70).

4) Eliminate memory leak when repeatedly executing the same query.

5) Eliminate segfault when attempting to reuse a REF cursor that has been closed.

6) Updated documentation.

13.2.10 Version 6.0.1 (August 2017)

1) Update to ODPI-C 2.0.1.

   - Ensure that queries registered via `Subscription.registerquery()` do not prevent the associated connection from being closed (ODPI-C issue 27).

   - Deprecated attribute `Subscription.id` as it was never intended to be exposed (ODPI-C issue 28). It will be dropped in version 6.1.

   - Add support for DML Returning statements that require dynamically allocated variable data (such as CLOBs being returned as strings).

2) Correct packaging of Python 2.7 UCS4 wheels on Linux (issue 64).

3) Updated documentation.

13.2.11 Version 6.0 (August 2017)

See *What's New* for a summary of the changes between cx_Oracle 5.3 and cx_Oracle 6.0.

1) Update to ODPI-C 2.0.

   - Prevent closing the connection when there are any open statements or LOBs and add new error "DPI-1054: connection cannot be closed when open statements or LOBs exist" when this situation is detected; this is needed to prevent crashes under certain conditions when statements or LOBs are being acted upon while at the same time (in another thread) a connection is being closed; it also prevents leaks of statements and LOBs when a connection is returned to a session pool.

   - On platforms other than Windows, if the regular method for loading the Oracle Client libraries fails, try using $ORACLE_HOME/lib/libclntsh.so (ODPI-C issue 20).

   - Use the environment variable DPI_DEBUG_LEVEL at runtime, not compile time.

   - Added support for DPI_DEBUG_LEVEL_ERRORS (reports errors and has the value 8) and DPI_DEBUG_LEVEL_SQL (reports prepared SQL statement text and has the value 16) in order to further improve the ability to debug issues.

   - Correct processing of `Cursor.scroll()` in some circumstances.

2) Delay initialization of the ODPI-C library until the first standalone connection or session pool is created so that manipulation of the environment variable NLS_LANG can be performed after the module has been imported; this also has the added benefit of reducing the number of errors that can take place when the module is imported.

3) Prevent binding of null values from generating the exception "ORA-24816: Expanded non LONG bind data supplied after actual LONG or LOB column" in certain circumstances (issue 50).

4) Added information on how to run the test suite (issue 33).

5) Documentation improvements.

13.2.12 Version 6.0 rc 2 (July 2017)

1) Update to ODPI-C rc 2.

   - Provide improved error message when OCI environment cannot be created, such as when the oraaccess.xml file cannot be processed properly.

   - On Windows, convert system message to Unicode first, then to UTF-8; otherwise, the error message returned could be in a mix of encodings (issue 40).

   - Corrected support for binding decimal values in object attribute values and collection element values.

   - Corrected support for binding PL/SQL boolean values to PL/SQL procedures with Oracle client 11.2.

2) Define exception classes on the connection object in addition to at module scope in order to simplify error handling in multi-connection environments, as specified in the Python DB API.

3) Ensure the correct encoding is used for setting variable values.

4) Corrected handling of CLOB/NCLOB when using different encodings.

5) Corrected handling of TIMESTAMP WITH TIME ZONE attributes on objects.

6) Ensure that the array position passed to var.getvalue() does not exceed the number of elements allocated in the array.

7) Reworked test suite and samples so that they are independent of each other and so that the SQL scripts used to create/drop schemas are easily adjusted to use different schema names, if desired.

8) Updated DB API test suite stub to support Python 3.

9) Added additional test cases and samples.

10) Documentation improvements.

13.2.13 Version 6.0 rc 1 (June 2017)

1) Update to ODPI-C rc 1.

2) The method `Cursor.setoutputsize()` no longer needs to do anything, since ODPI-C automatically manages buffer sizes of LONG and LONG RAW columns.

3) Handle case when both precision and scale are zero, as occurs when retrieving numeric expressions (issue 34).

4) OCI requires that both encoding and nencoding have values or that both encoding and encoding do not have values. These parameters are used in functions `cx_Oracle.connect()` and `cx_Oracle. SessionPool()`. The missing value is set to its default value if one of the values is set and the other is not (issue 36).

5) Permit use of both string and unicode for Python 2.7 for creating session pools and for changing passwords (issue 23).

6) Corrected handling of BFILE LOBs.

7) Add script for dropping test schemas.

8) Documentation improvements.

13.2.14 Version 6.0 beta 2 (May 2017)

1) Added support for getting/setting attributes of objects or element values in collections that contain LOBs, BI-NARY_FLOAT values, BINARY_DOUBLE values and NCHAR and NVARCHAR2 values. The error message for any types that are not supported has been improved as well.

2) Enable temporary LOB caching in order to avoid disk I/O as suggested.

3) Added support for setting the debug level in ODPI-C, if desirable, by setting environment variable DPI_DEBUG_LEVEL prior to building cx_Oracle.

4) Correct processing of strings in `Cursor.executemany()` when a larger string is found after a shorter string in the list of data bound to the statement.

5) Correct handling of long Python integers that cannot fit inside a 64-bit C integer (issue 18).

6) Correct creation of pool using external authentication.

7) Handle edge case when an odd number of zeroes trail the decimal point in a value that is effectively zero (issue 22).

8) Prevent segfault under load when the attempt to create an error fails.

9) Eliminate resource leak when a standalone connection or pool is freed.

10) Correct typo.

11) Correct handling of REF cursors when the array size is manipulated.

12) Prevent attempts from binding the cursor being executed to itself.

13) Correct reference count handling of parameters when creating a cursor.

14) Correct determination of the names of the bind variables in prepared SQL statements (which behaves a little differently from PL/SQL statements).

13.2.15 Version 6.0 beta 1 (April 2017)

1) Simplify building cx_Oracle considerably by use of ODPI-C. This means that cx_Oracle can now be built without Oracle Client header files or libraries and that at runtime cx_Oracle can adapt to Oracle Client 11.2, 12.1 or 12.2 libraries without needing to be rebuilt. This also means that wheels can now be produced and installed via pip.

2) Added attribute `SessionPool.stmtcachesize` to support getting and setting the default statement cache size for connections in the pool.

3) Added attribute `Connection.dbop` to support setting the database operation that is to be monitored.

4) Added attribute `Connection.handle` to facilitate testing the creation of a connection using a OCI service context handle.

5) Added parameters tag and matchanytag to the `cx_Oracle.connect()` and `SessionPool.acquire()` methods and added parameters tag and retag to the `SessionPool.release()` method in order to support session tagging.

6) Added parameter edition to the `cx_Oracle.SessionPool()` method.

7) Added support for universal rowids.

8) Added support for DML Returning of multiple rows.

9) Added attributes `Variable.actualElements` and `Variable.values` to variables.

10) Added parameters region, sharding_key and super_sharding_key to the `cx_Oracle.makedsn()` method to support connecting to a sharded database (new in Oracle Database 12.2).

11) Added support for smallint and float data types in Oracle objects, as requested.

12) An exception is no longer raised when a collection is empty for methods `Object.first()` and `Object.last()`. Instead, the value None is returned to be consistent with the methods `Object.next()` and `Object.prev()`.

13) If the environment variables NLS_LANG and NLS_NCHAR are being used, they must be set before the module is imported. Using the encoding and nencoding parameters to the `cx_Oracle.connect()` and `cx_Oracle.SessionPool()` methods is a simpler alternative to setting these environment variables.

14) Removed restriction on fetching LOBs across round trips to the database (eliminates error "LOB variable no longer valid after subsequent fetch").

15) Removed requirement for specifying a maximum size when fetching LONG or LONG raw columns. This also allows CLOB, NCLOB, BLOB and BFILE columns to be fetched as strings or bytes without needing to specify a maximum size.

16) Dropped deprecated parameter twophase from the `cx_Oracle.connect()` method. Applications should set the `Connection.internal_name` and `Connection.external_name` attributes instead to a value appropriate to the application.

17) Dropped deprecated parameters action, module and clientinfo from the `cx_Oracle.connect()` method. The appcontext parameter should be used instead as shown in this sample.

18) Dropped deprecated attribute numbersAsString from *cursor objects*. Use an output type handler instead as shown in this sample.

19) Dropped deprecated attributes cqqos and rowids from *subscription objects*. Use the qos attribute instead as shown in this sample.

20) Dropped deprecated parameters cqqos and rowids from the `Connection.subscribe()` method. Use the qos parameter instead as shown in this sample.

13.3 5.x releases

13.3.1 Version 5.3 (March 2017)

1) Added support for Python 3.6.

2) Dropped support for Python versions earlier than 2.6.

3) Dropped support for Oracle clients earlier than 11.2.

4) Added support for `fetching implicit results` (available in Oracle 12.1)

5) Added support for `Transaction Guard` (available in Oracle 12.1).

6) Added support for setting the `maximum lifetime` of pool connections (available in Oracle 12.1).

7) Added support for large row counts (larger than 2 ** 32, available in Oracle 12.1)

8) Added support for `advanced queuing`.

9) Added support for `scrollable cursors`.

10) Added support for `edition based redefinition`.

11) Added support for `creating`, modifying and binding user defined types and collections.

12) Added support for creating, modifying and binding PL/SQL records and collections (available in Oracle 12.1).

13) Added support for binding *native integers*.

14) Enabled statement caching.

15) Removed deprecated variable attributes maxlength and allocelems.

16) Corrected support for setting the encoding and nencoding parameters when *creating a connection* and added support for setting these when creating a session pool. These can now be used instead of setting the environment variables NLS_LANG and NLS_NCHAR.

17) Use None instead of 0 for items in the *Cursor.description* attribute that do not have any validity.

18) Changed driver name to match informal driver name standard used by Oracle for other drivers.

19) Add check for maximum of 10,000 parameters when calling a stored procedure or function in order to prevent a possible improper memory access from taking place.

20) Removed -mno-cygwin compile flag since it is no longer used in newer versions of the gcc compiler for Cygwin.

21) Simplified test suite by combining Python 2 and 3 scripts into one script and separated out 12.1 features into a single script.

22) Updated samples to use code that works on both Python 2 and 3

23) Added support for pickling/unpickling error objects (Issue #23)

24) Dropped support for callbacks on OCI functions.

25) Removed deprecated types UNICODE, FIXED_UNICODE and LONG_UNICODE (use NCHAR, FIXED_NCHAR and LONG_NCHAR instead).

26) Increased default array size to 100 (from 50) to match other drivers.

27) Added support for setting the *internal_name* and *external_name* on the connection directly. The use of the twophase parameter is now deprecated. Applications should set the internal_name and external_name attributes directly to a value appropriate to the application.

28) Added support for using application context when *creating a connection*. This should be used in preference to the module, action and clientinfo parameters which are now deprecated.

29) Reworked database change notification and continuous query notification to more closely align with the PL/SQL implementation and prepare for sending notifications for AQ messages. The following changes were made:

   • added constant *SUBSCR_QOS_BEST_EFFORT* to replace deprecated constant SUB-SCR_CQ_QOS_BEST_EFFORT

   • added constant *SUBSCR_QOS_QUERY* to replace deprecated constant SUBSCR_CQ_QOS_QUERY

   • added constant *SUBSCR_QOS_DEREG_NFY* to replace deprecated constant SUB-SCR_QOS_PURGE_ON_NTFN

   • added constant *SUBSCR_QOS_ROWIDS* to replace parameter rowids for method *Connection.subscribe()*

   • deprecated parameter cqqos for method *Connection.subscribe()*. The qos parameter should be used instead.

   • dropped constants SUBSCR_CQ_QOS_CLQRYCACHE, SUBSCR_QOS_HAREG, SUB-SCR_QOS_MULTICBK, SUBSCR_QOS_PAYLOAD, SUBSCR_QOS_REPLICATE, and SUB-SCR_QOS_SECURE since they were never actually used

30) Deprecated use of the numbersAsStrings attribute on cursors. An output type handler should be used instead.

13.3.2 Version 5.2.1 (January 2016)

1) Added support for Python 3.5.

2) Removed password attribute from connection and session pool objects in order to promote best security practices (if stored in RAM in cleartext it can be read in process dumps, for example). For those who would like to retain this feature, a subclass of Connection could be used to store the password.

3) Added optional parameter externalauth to SessionPool() which enables wallet based or other external authentication mechanisms to be used.

4) Use the national character set encoding when required (when char set form is SQLCS_NCHAR); otherwise, the wrong encoding would be used if the environment variable NLS_NCHAR is set.

5) Added support for binding boolean values to PL/SQL blocks and stored procedures (available in Oracle 12.1).

13.3.3 Version 5.2 (June 2015)

1) Added support for strings up to 32k characters (new in Oracle 12c).

2) Added support for getting array DML row counts (new in Oracle 12c).

3) Added support for fetching batch errors.

4) Added support for LOB values larger than 4 GB.

5) Added support for connections as SYSASM.

6) Added support for building without any configuration changes to the machine when using instant client RPMs on Linux.

7) Added types NCHAR, FIXED_NCHAR and LONG_NCHAR to replace the types UNICODE, FIXED_UNICODE and LONG_UNICODE (which are now deprecated). These types are available in Python 3 as well so they can be used to specify the use of NCHAR type fields when binding or using setinputsizes().

8) Fixed binding of booleans in Python 3.x.

9) Test suite now sets NLS_LANG if not already set.

10) Enhanced documentation for connection.action attribute and added note on cursor.parse() method to make clear that DDL statements are executed when parsed.

11) Removed remaining remnants of support Oracle 9i.

12) Added __version__ attribute to conform with PEP 396.

13) Ensure that sessions are released to the pool when calling connection.close() (Issue #2)

14) Fixed handling of datetime intervals (Issue #7)

13.3.4 Version 5.1.3 (May 2014)

1) Added support for Oracle 12c.

2) Added support for Python 3.4.

3) Added support for query result set change notification. Thanks to Glen Walker for the patch.

4) Ensure that in Python 3.x that NCHAR and NVARCHAR2 and NCLOB columns are retrieved properly without conversion issues. Thanks to Joakim Andersson for pointing out the issue and the possible solution.

5) Fix bug when an exception is caught and then another exception is raised while handling that exception in Python 3.x. Thanks to Boris Dzuba for pointing out the issue and providing a test case.

6) Enhance performance returning integers between 10 and 18 digits on 64-bit platforms that support it. Thanks for Shai Berger for the initial patch.

7) Fixed two memory leaks.

8) Fix to stop current_schema from throwing a MemoryError on 64-bit platforms on occasion. Thanks to Andrew Horton for the fix.

9) Class name of cursors changed to real name cx_Oracle.Cursor.

13.3.5 Version 5.1.2 (July 2012)

1) Added support for LONG_UNICODE which is a type used to handle long unicode strings. These are not explicitly supported in Oracle but can be used to bind to NCLOB, for example, without getting the error "unimplemented or unreasonable conversion requested".

2) Set the row number in a cursor when executing PL/SQL blocks as requested by Robert Ritchie.

3) Added support for setting the module, action and client_info attributes during connection so that logon triggers will see the supplied values, as requested by Rodney Barnett.

13.3.6 Version 5.1.1 (October 2011)

1) Simplify management of threads for callbacks performed by database change notification and eliminate a crash that occurred under high load in certain situations. Thanks to Calvin S. for noting the issue and suggesting a solution and testing the patch.

2) Force server detach on close so that the connection is completely closed and not just the session as before.

3) Force use of OCI_UTF16ID for NCLOBs as using the default character set would result in ORA-03127 with Oracle 11.2.0.2 and UTF8 character set.

4) Avoid attempting to clear temporary LOBs a second time when destroying the variable as in certain situations this results in spurious errors.

5) Added additional parameter service_name to makedsn() which can be used to use the service_name rather than the SID in the DSN string that is generated.

6) Fix cursor description in test suite to take into account the number of bytes per character.

7) Added tests for NCLOBS to the test suite.

8) Removed redundant code in setup.py for calculating the library path.

13.3.7 Version 5.1 (March 2011)

1) Remove support for UNICODE mode and permit Unicode to be passed through in everywhere a string may be passed in. This means that strings will be passed through to Oracle using the value of the NLS_LANG environment variable in Python 3.x as well. Doing this eliminated a bunch of problems that were discovered by using UNICODE mode and also removed an unnecessary restriction in Python 2.x that Unicode could not be used in connect strings or SQL statements, for example.

2) Added support for creating an empty object variable via a named type, the first step to adding full object support.

3) Added support for Python 3.2.

4) Account for lib64 used on x86_64 systems. Thanks to Alex Wood for supplying the patch.

5) Clear up potential problems when calling cursor.close() ahead of the cursor being freed by going out of scope.

6) Avoid compilation difficulties on AIX5 as OCIPing does not appear to be available on that platform under Oracle 10g Release 2. Thanks to Pierre-Yves Fontaniere for the patch.

7) Free temporary LOBs prior to each fetch in order to avoid leaking them. Thanks to Uwe Hoffmann for the initial patch.

13.3.8 Version 5.0.4 (July 2010)

1) Added support for Python 2.7.

2) Added support for new parameter (port) for subscription() call which allows the client to specify the listening port for callback notifications from the database server. Thanks to Geoffrey Weber for the initial patch.

3) Fixed compilation under Oracle 9i.

4) Fixed a few error messages.

13.3.9 Version 5.0.3 (February 2010)

1) Added support for 64-bit Windows.

2) Added support for Python 3.1 and dropped support for Python 3.0.

3) Added support for keyword parameters in cursor.callproc() and cursor.callfunc().

4) Added documentation for the UNICODE and FIXED_UNICODE variable types.

5) Added extra link arguments required for Mac OS X as suggested by Jason Woodward.

6) Added additional error codes to the list of error codes that raise OperationalError rather than DatabaseError.

7) Fixed calculation of display size for strings with national database character sets that are not the default AL16UTF16.

8) Moved the resetting of the setinputsizes flag before the binding takes place so that if an error takes place and a new statement is prepared subsequently, spurious errors will not occur.

9) Fixed compilation with Oracle 10g Release 1.

10) Tweaked documentation based on feedback from a number of people.

11) Added support for running the test suite using "python setup.py test"

12) Added support for setting the CLIENT_IDENTIFIER value in the v$session table for connections.

13) Added exception when attempting to call executemany() with arrays which is not supported by the OCI.

14) Fixed bug when converting from decimal would result in OCI-22062 because the locale decimal point was not a period. Thanks to Amaury Forgeot d'Arc for the solution to this problem.

13.3.10 Version 5.0.2 (May 2009)

1) Fix creation of temporary NCLOB values and the writing of NCLOB values in non Unicode mode.

2) Re-enabled parsing of non select statements as requested by Roy Terrill.

3) Implemented a parse error offset as requested by Catherine Devlin.

4) Removed lib subdirectory when forcing RPATH now that the library directory is being calculated exactly in setup.py.

5) Added an additional cast in order to support compiling by Microsoft Visual C++ 2008 as requested by Marco de Paoli.

6) Added additional include directory to setup.py in order to support compiling by Microsoft Visual Studio was requested by Jason Coombs.

7) Fixed a few documentation issues.

13.3.11 Version 5.0.1 (February 2009)

1) Added support for database change notification available in Oracle 10g Release 2 and higher.

2) Fix bug where NCLOB data would be corrupted upon retrieval (non Unicode mode) or would generate exception ORA-24806 (LOB form mismatch). Oracle insists upon differentiating between CLOB and NCLOB no matter which character set is being used for retrieval.

3) Add new attributes size, bufferSize and numElements to variable objects, deprecating allocelems (replaced by numElements) and maxlength (replaced by bufferSize)

4) Avoid increasing memory allocation for strings when using variable width character sets and increasing the number of elements in a variable during executemany().

5) Tweaked code in order to ensure that cx_Oracle can compile with Python 3.0.1.

13.3.12 Version 5.0 (December 2008)

1) Added support for Python 3.0 with much help from Amaury Forgeot d'Arc.

2) Removed support for Python 2.3 and Oracle 8i.

3) Added support for full unicode mode in Python 2.x where all strings are passed in and returned as unicode (module must be built in this mode) rather than encoded strings

4) nchar and nvarchar columns now return unicode instead of encoded strings

5) Added support for an output type handler and/or an input type handler to be specified at the connection and cursor levels.

6) Added support for specifying both input and output converters for variables

7) Added support for specifying the array size of variables that are created using the cursor.var() method

8) Added support for events mode and database resident connection pooling (DRCP) in Oracle 11g.

9) Added support for changing the password during construction of a new connection object as well as after the connection object has been created

10) Added support for the interval day to second data type in Oracle, represented as datetime.timedelta objects in Python.

11) Added support for getting and setting the current_schema attribute for a session

12) Added support for proxy authentication in session pools as requested by Michael Wegrzynek (and thanks for the initial patch as well).

13) Modified connection.prepare() to return a boolean indicating if a transaction was actually prepared in order to avoid the error ORA-24756 (transaction does not exist).

14) Raise a cx_Oracle.Error instance rather than a string for column truncation errors as requested by Helge Tesdal.

15) Fixed handling of environment handles in session pools in order to allow session pools to fetch objects without exceptions taking place.

13.4 Older releases

13.4.1 Version 4.4.1 (October 2008)

1) Make the bind variables and fetch variables accessible although they need to be treated carefully since they are used internally; support added for forward compatibility with version 5.x.

2) Include the "cannot insert null value" in the list of errors that are treated as integrity errors as requested by Matt Boersma.

3) Use a cx_Oracle.Error instance rather than a string to hold the error when truncation (ORA-1406) takes place as requested by Helge Tesdal.

4) Added support for fixed char, old style varchar and timestamp attribute values in objects.

5) Tweaked setup.py to check for the Oracle version up front rather than during the build in order to produce more meaningful errors and simplify the code.

6) In setup.py added proper detection for the instant client on Mac OS X as recommended by Martijn Pieters.

7) In setup.py, avoided resetting the extraLinkArgs on Mac OS X as doing so prevents simple modification where desired as expressed by Christian Zagrodnick.

8) Added documentation on exception handling as requested by Andreas Mock, who also graciously provided an initial patch.

9) Modified documentation indicating that the password attribute on connection objects can be written.

10) Added documentation warning that parameters not passed in during subsequent executions of a statement will retain their original values as requested by Harald Armin Massa.

11) Added comments indicating that an Oracle client is required since so many people find this surprising.

12) Removed all references to Oracle 8i from the documentation and version 5.x will eliminate all vestiges of support for this version of the Oracle client.

13) Added additional link arguments for Cygwin as requested by Rob Gillen.

13.4.2 Version 4.4 (June 2008)

1) Fix setup.py to handle the Oracle instant client and Oracle XE on both Linux and Windows as pointed out by many. Thanks also to the many people who also provided patches.

2) Set the default array size to 50 instead of 1 as the DB API suggests because the performance difference is so drastic and many people have recommended that the default be changed.

3) Added Py_BEGIN_ALLOW_THREADS and Py_END_ALLOW_THREADS around each blocking call for LOBs as requested by Jason Conroy who also provided an initial patch and performed a number of tests that demonstrate the new code is much more responsive.

4) Add support for acquiring cursor.description after a parse.

5) Defer type assignment when performing executemany() until the last possible moment if the value being bound in is null as suggested by Dragos Dociu.

6) When dropping a connection from the pool, ignore any errors that occur during the rollback; unfortunately, Oracle decides to commit data even when dropping a connection from the pool instead of rolling it back so the attempt still has to be made.

7) Added support for setting CLIENT_DRIVER in V$SESSION_CONNECT_INFO in Oracle 11g and higher.

8) Use cx_Oracle.InterfaceError rather than the builtin RuntimeError when unable to create the Oracle environment object as requested by Luke Mewburn since the error is specific to Oracle and someone attempting to catch any exception cannot simply use cx_Oracle.Error.

9) Translated some error codes to OperationalError as requested by Matthew Harriger; translated if/elseif/else logic to switch statement to make it more readable and to allow for additional translation if desired.

10) Transformed documentation to new format using restructured text. Thanks to Waldemar Osuch for contributing the initial draft of the new documentation.

11) Allow the password to be overwritten by a new value as requested by Alex VanderWoude; this value is retained as a convenience to the user and not used by anything in the module; if changed externally it may be convenient to keep this copy up to date.

12) Cygwin is on Windows so should be treated in the same way as noted by Matthew Cahn.

13) Add support for using setuptools if so desired as requested by Shreya Bhatt.

14) Specify that the version of Oracle 10 that is now primarily used is 10.2, not 10.1.

13.4.3 Version 4.3.3 (October 2007)

1) Added method ping() on connections which can be used to test whether or not a connection is still active (available in Oracle 10g R2).

2) Added method cx_Oracle.clientversion() which returns a 5-tuple giving the version of the client that is in use (available in Oracle 10g R2).

3) Added methods startup() and shutdown() on connections which can be used to startup and shutdown databases (available in Oracle 10g R2).

4) Added support for Oracle 11g.

5) Added samples directory which contains a handful of scripts containing sample code for more advanced techniques. More will follow in future releases.

6) Prevent error "ORA-24333: zero iteration count" when calling executemany() with zero rows as requested by Andreas Mock.

7) Added methods __enter__() and __exit__() on connections to support using connections as context managers in Python 2.5 and higher. The context managed is the transaction state. Upon exit the transaction is either rolled back or committed depending on whether an exception took place or not.

8) Make the search for the lib32 and lib64 directories automatic for all platforms.

9) Tweak the setup configuration script to include all of the metadata and allow for building the module within another setup configuration script

10) Include the Oracle version in addition to the Python version in the build directories that are created and in the names of the binary packages that are created.

11) Remove unnecessary dependency on win32api to build module on Windows.

13.4.4 Version 4.3.2 (August 2007)

1) Added methods open(), close(), isopen() and getchunksize() in order to improve performance of reading/writing LOB values in chunks.

2) Fixed support for native doubles and floats in Oracle 10g; added new type NATIVE_FLOAT to allow specification of a variable of that specific type where desired. Thanks to D.R. Boxhoorn for pointing out the fact that this was not working properly when the arraysize was anything other than 1.

3) When calling connection.begin(), only create a new transaction handle if one is not already associated with the connection. Thanks to Andreas Mock for discovering this and for Amaury Forgeot d'Arc for diagnosing the problem and pointing the way to a solution.

4) Added attribute cursor.rowfactory which allows a method to be called for each row that is returned; this is about 20% faster than calling the method in Python using the idiom [method(*r) for r in cursor].

5) Attempt to locate an Oracle installation by looking at the PATH if the environment variable ORACLE_HOME is not set; this is of primary use on Windows where this variable should not normally be set.

6) Added support for autocommit mode as requested by Ian Kelly.

7) Added support for connection.stmtcachesize which allows for both reading and writing the size of the statement cache size. This parameter can make a huge difference with the length of time taken to prepare statements. Added support for setting the statement tag when preparing a statement. Both of these were requested by Bjorn Sandberg who also provided an initial patch.

8) When copying the value of a variable, copy the return code as well.

13.4.5 Version 4.3.1 (April 2007)

1) Ensure that if the client buffer size exceeds 4000 bytes that the server buffer size does not as strings may only contain 4000 bytes; this allows handling of multibyte character sets on the server as well as the client.

2) Added support for using buffer objects to populate binary data and made the Binary() constructor the buffer type as requested by Ken Mason.

3) Fix potential crash when using full optimization with some compilers. Thanks to Aris Motas for noticing this and providing the initial patch and to Amaury Forgeot d'Arc for providing an even simpler solution.

4) Pass the correct charset form in to the write call in order to support writing to national character set LOB values properly. Thanks to Ian Kelly for noticing this discrepancy.

13.4.6 Version 4.3 (March 2007)

1) Added preliminary support for fetching Oracle objects (SQL types) as requested by Kristof Beyls (who kindly provided an initial patch). Additional work needs to be done to support binding and updating objects but the basic structure is now in place.

2) Added connection.maxBytesPerCharacter which indicates the maximum number of bytes each character can use; use this value to also determine the size of local buffers in order to handle discrepancies between the client character set and the server character set. Thanks to Andreas Mock for providing the initial patch and working with me to resolve this issue.

3) Added support for querying native floats in Oracle 10g as requested by Danny Boxhoorn.

4) Add support for temporary LOB variables created via PL/SQL instead of only directly by cx_Oracle; thanks to Henning von Bargen for discovering this problem.

5) Added support for specifying variable types using the builtin types int, float, str and datetime.date which allows for finer control of what type of Python object is returned from cursor.callfunc() for example.

6) Added support for passing booleans to callproc() and callfunc() as requested by Anana Aiyer.

7) Fixed support for 64-bit environments in Python 2.5.

8) Thanks to Filip Ballegeer and a number of his co-workers, an intermittent crash was tracked down; specifically, if a connection is closed, then the call to OCIStmtRelease() will free memory twice. Preventing the call when the connection is closed solves the problem.

13.4.7 Version 4.2.1 (September 2006)

1) Added additional type (NCLOB) to handle CLOBs that use the national character set as requested by Chris Dunscombe.

2) Added support for returning cursors from functions as requested by Daniel Steinmann.

3) Added support for getting/setting the "get" mode on session pools as requested by Anand Aiyer.

4) Added support for binding subclassed cursors.

5) Fixed binding of decimal objects with absolute values less than 0.1.

13.4.8 Version 4.2 (July 2006)

1) Added support for parsing an Oracle statement as requested by Patrick Blackwill.

2) Added support for BFILEs at the request of Matthew Cahn.

3) Added support for binding decimal.Decimal objects to cursors.

4) Added support for reading from NCLOBs as requested by Chris Dunscombe.

5) Added connection attributes encoding and nencoding which return the IANA character set name for the character set and national character set in use by the client.

6) Rework module initialization to use the techniques recommended by the Python documentation as one user was experiencing random segfaults due to the use of the module dictionary after the initialization was complete.

7) Removed support for the OPT_Threading attribute. Use the threaded keyword when creating connections and session pools instead.

8) Removed support for the OPT_NumbersAsStrings attribute. Use the numbersAsStrings attribute on cursors instead.

9) Use type long rather than type int in order to support long integers on 64-bit machines as reported by Uwe Hoffmann.

10) Add cursor attribute "bindarraysize" which is defaulted to 1 and is used to determine the size of the arrays created for bind variables.

11) Added repr() methods to provide something a little more useful than the standard type name and memory address.

12) Added keyword parameter support to the functions that imply such in the documentation as requested by Harald Armin Massa.

13) Treat an empty dictionary passed through to cursor.execute() as keyword parameters the same as if no keyword parameters were specified at all, as requested by Fabien Grumelard.

14) Fixed memory leak when a LOB read would fail.

---

15) Set the LDFLAGS value in the environment rather than directly in the setup.py file in order to satisfy those who wish to enable the use of debugging symbols.

16) Use __DATE__ and __TIME__ to determine the date and time of the build rather than passing it through directly.

17) Use Oracle types and add casts to reduce warnings as requested by Amaury Forgeot d'Arc.

18) Fixed typo in error message.

13.4.9 Version 4.1.2 (December 2005)

1) Restore support of Oracle 9i features when using the Oracle 10g client.

13.4.10 Version 4.1.1 (December 2005)

1) Add support for dropping a connection from a session pool.

2) Add support for write only attributes "module", "action" and "clientinfo" which work only in Oracle 10g as requested by Egor Starostin.

3) Add support for pickling database errors.

4) Use the previously created bind variable as a template if available when creating a new variable of a larger size. Thanks to Ted Skolnick for the initial patch.

5) Fixed tests to work properly in the Python 2.4 environment where dates and timestamps are different Python types. Thanks to Henning von Bargen for pointing this out.

6) Added additional directories to search for include files and libraries in order to better support the Oracle 10g instant client.

7) Set the internal fetch number to 0 in order to satisfy very picky source analysis tools as requested by Amaury Fogeot d'Arc.

8) Improve the documentation for building and installing the module from source as some people are unaware of the standard methods for building Python modules using distutils.

9) Added note in the documentation indicating that the arraysize attribute can drastically affect performance of queries since this seems to be a common misunderstanding of first time users of cx_Oracle.

10) Add a comment indicating that on HP-UX Itanium with Oracle 10g the library ttsh10 must also be linked against. Thanks to Bernard Delmee for the information.

13.4.11 Version 4.1 (January 2005)

1) Fixed bug where subclasses of Cursor do not pass the connection in the constructor causing a segfault.

2) DDL statements must be reparsed before execution as noted by Mihai Ibanescu.

3) Add support for setting input sizes by position.

4) Fixed problem with catching an exception during execute and then still attempting to perform a fetch afterwards as noted by Leith Parkin.

5) Rename the types so that they can be pickled and unpickled. Thanks to Harri Pasanen for pointing out the problem.

6) Handle invalid NLS_LANG setting properly (Oracle seems to like to provide a handle back even though it is invalid) and determine the number of bytes per character in order to allow for proper support in the future of multibyte and variable width character sets.

---

7) Remove date checking from the native case since Python already checks that dates are valid; enhance error message when invalid dates are encountered so that additional processing can be done.

8) Fix bug executing SQL using numeric parameter names with predefined variables (such as what takes place when calling stored procedures with out parameters).

9) Add support for reading CLOB values using multibyte or variable length character sets.

13.4.12 Version 4.1 beta 1 (September 2004)

1) Added support for Python 2.4. In Python 2.4, the datetime module is used for both binding and fetching of date and timestamp data. In Python 2.3, objects from the datetime module can be bound but the internal datetime objects will be returned from queries.

2) Added pickling support for LOB and datetime data.

3) Fully qualified the table name that was missing in an alter table statement in the setup test script as noted by Marc Gehling.

4) Added a section allowing for the setting of the RPATH linker directive in setup.py as requested by Iustin Pop.

5) Added code to raise a programming error exception when an attempt is made to access a LOB locator variable in a subsequent fetch.

6) The username, password and dsn (tnsentry) are stored on the connection object when specified, regardless of whether or not a standard connection takes place.

7) Added additional module level constant called "LOB" as requested by Joseph Canedo.

8) Changed exception type to IntegrityError for constraint violations as requested by Joseph Canedo.

9) If scale and precision are not specified, an attempt is made to return a long integer as requested by Joseph Canedo.

10) Added workaround for Oracle bug which returns an invalid handle when the prepare call fails. Thanks to alantam@hsbc.com for providing the code that demonstrated the problem.

11) The cursor method arrayvar() will now accept the actual list so that it is not necessary to call cursor.arrayvar() followed immediately by var.setvalue().

12) Fixed bug where attempts to execute the statement "None" with bind variables would cause a segmentation fault.

13) Added support for binding by position (paramstyle = "numeric").

14) Removed memory leak created by calls to OCIParamGet() which were not mirrored by calls to OCIDescriptor-Free(). Thanks to Mihai Ibanescu for pointing this out and providing a patch.

15) Added support for calling cursor.executemany() with statement None implying that the previously prepared statement ought to be executed. Thanks to Mihai Ibanescu for providing a patch.

16) Added support for rebinding variables when a subsequent call to cursor.executemany() uses a different number of rows. Thanks to Mihai Ibanescu for supplying a patch.

17) The microseconds are now displayed in datetime variables when nonzero similar to method used in the datetime module.

18) Added support for binary_float and binary_double columns in Oracle 10g.

13.4.13 Version 4.0.1 (February 2004)

1) Fixed bugs on 64-bit platforms that caused segmentation faults and bus errors in session pooling and determining the bind variables associated with a statement.

2) Modified test suite so that 64-bit platforms are tested properly.

3) Added missing commit statements in the test setup scripts. Thanks to Keith Lyon for pointing this out.

4) Fix setup.py for Cygwin environments. Thanks to Doug Henderson for providing the necessary fix.

5) Added support for compiling cx_Oracle without thread support. Thanks to Andre Reitz for pointing this out.

6) Added support for a new keyword parameter called threaded on connections and session pools. This parameter defaults to False and indicates whether threaded mode ought to be used. It replaces the module level attribute OPT_Threading although examining the attribute will be retained until the next release at least.

7) Added support for a new keyword parameter called twophase on connections. This parameter defaults to False and indicates whether support for two phase (distributed or global) transactions ought to be present. Note that support for distributed transactions is buggy when crossing major version boundaries (Oracle 8i to Oracle 9i for example).

8) Ensure that the rowcount attribute is set properly when an exception is raised during execution. Thanks to Gary Aviv for pointing out this problem and its solution.

13.4.14 Version 4.0 (December 2003)

1) Added support for subclassing connections, cursors and session pools. The changes involved made it necessary to drop support for Python 2.1 and earlier although a branch exists in CVS to allow for support of Python 2.1 and earlier if needed.

2) Connections and session pools can now be created with keyword parameters, not just sequential parameters.

3) Queries now return integers whenever possible and long integers if the number will overflow a simple integer. Floats are only returned when it is known that the number is a floating point number or the integer conversion fails.

4) Added initial support for user callbacks on OCI functions. See the documentation for more details.

5) Add support for retrieving the bind variable names associated with a cursor with a new method bindnames().

6) Add support for temporary LOB variables. This means that setinputsizes() can be used with the values CLOB and BLOB to create these temporary LOB variables and allow for the equivalent of empty_clob() and empty_blob() since otherwise Oracle will treat empty strings as NULL values.

7) Automatically switch to long strings when the data size exceeds the maximum string size that Oracle allows (4000 characters) and raise an error if an attempt is made to set a string variable to a size that it does not support. This avoids truncation errors as reported by Jon Franz.

8) Add support for global (distributed) transactions and two phase commit.

9) Force the NLS settings for the session so that test tables are populated correctly in all circumstances; problems were noted by Ralf Braun and Allan Poulsen.

10) Display error messages using the environment handle when the error handle has not yet been created; this provides better error messages during this rather rare situation.

11) Removed memory leak in callproc() that was reported by Todd Whiteman.

12) Make consistent the calls to manipulate memory; otherwise segfaults can occur when the pymalloc option is used, as reported by Matt Hoskins.

13) Force a rollback when a session is released back to the session pool. Apparently the connections are not as stateless as Oracle's documentation suggests and this makes the logic consistent with normal connections.

14) Removed module method attach(). This can be replaced with a call to Connection(handle=) if needed.

13.4.15 Version 3.1 (August 2003)

1) Added support for connecting with SYSDBA and SYSOPER access which is needed for connecting as sys in Oracle 9i.

2) Only check the dictionary size if the variable is not NULL; otherwise, an error takes place which is not caught or cleared; this eliminates a spurious "Objects/dictobject.c:1258: bad argument to internal function" in Python 2.3.

3) Add support for session pooling. This is only support for Oracle 9i but is amazingly fast – about 100 times faster than connecting.

4) Add support for statement caching when pooling sessions, this reduces the parse time considerably. Unfortunately, the Oracle OCI does not allow this to be easily turned on for normal sessions.

5) Add method trim() on CLOB and BLOB variables for trimming the size.

6) Add support for externally identified users; to use this feature leave the username and password fields empty when connecting.

7) Add method cancel() on connection objects to cancel long running queries. Note that this only works on non-Windows platforms.

8) Add method callfunc() on cursor objects to allow calling a function without using an anonymous PL/SQL block.

9) Added documentation on objects that were not documented. At this point all objects, methods and constants in cx_Oracle have been documented.

10) Added support for timestamp columns in Oracle 9i.

11) Added module level method makedsn() which creates a data source name given the host, port and SID.

12) Added constant "buildtime" which is the time when the module was built as an additional means of identifying the build that is in use.

13) Binding a value that is incompatible to the previous value that was bound (data types do not match or array size is larger) will now result in a new bind taking place. This is more consistent with the DB API although it does imply a performance penalty when used.

13.4.16 Version 3.0a (June 2003)

1) Fixed bug where zero length PL/SQL arrays were being mishandled

2) Fixed support for the data type "float" in Oracle; added one to the display size to allow for the sign of the number, if necessary; changed the display size of unconstrained numbers to 127, which is the largest number that Oracle can handle

3) Added support for retrieving the description of a bound cursor before fetching it

4) Fixed a couple of build issues on Mac OS X, AIX and Solaris (64-bit)

5) Modified documentation slightly based on comments from several people

6) Included files in MANIFEST that are needed to generate the binaries

7) Modified test suite to work within the test environment at Computronix as well as within the packages that are distributed

13.4.17 Version 3.0 (March 2003)

1) Removed support for connection to Oracle7 databases; it is entirely possible that it will still work but I no longer have any way of testing and Oracle has dropped any meaningful support for Oracle7 anyway

2) Fetching of strings is now done with predefined memory areas rather than dynamic memory areas; dynamic fetching of strings was causing problems with Oracle 9i in some instances and databases using a different character set other than US ASCII

3) Fixed bug where segfault would occur if the '/' character preceded the '@' character in a connect string

4) Added two new cursor methods var() and arrayvar() in order to eliminate the need for setinputsizes() when defining PL/SQL arrays and as a generic method of acquiring bind variables directly when needed

5) Fixed support for binding cursors and added support for fetching cursors (these are known as ref cursors in PL/SQL).

6) Eliminated discrepancy between the array size used internally and the array size specified by the interface user; this was done earlier to avoid bus errors on 64-bit platforms but another way has been found to get around that issue and a number of people were getting confused because of the discrepancy

7) Added support for the attribute "connection" on cursors, an optional DB API extension

8) Added support for passing a dictionary as the second parameter for the cursor.execute() method in order to comply with the DB API more closely; the method of passing parameters with keyword parameters is still supported and is in fact preferred

9) Added support for the attribute "statement" on cursors which is a reference to the last SQL statement prepared or executed

10) Added support for passing any sequence to callproc() rather than just lists as before

11) Fixed bug where segfault would occur if the array size was changed after the cursor was executed but before it was fetched

12) Ignore array size when performing executemany() and use the length of the list of parameters instead

13) Rollback when connection is closed or destroyed to follow DB API rather than use the Oracle default (which is commit)

14) Added check for array size too large causing an integer overflow

15) Added support for iterators for Python 2.2 and above

16) Added test suite based on PyUnitTest

17) Added documentation in HTML format similar to the documentation for the core Python library

13.4.18 Version 2.5a (August 2002)

1) Fix problem with Oracle 9i and retrieving strings; it seems that Oracle 9i uses the correct method for dynamic callback but Oracle 8i will not work with that method so an #ifdef was added to check for the existence of an Oracle 9i feature; thanks to Paul Denize for discovering this problem

13.4.19 Version 2.5 (July 2002)

1) Added flag OPT_NoOracle7 which, if set, assumes that connections are being made to Oracle8 or higher databases; this allows for eliminating the overhead in performing this check at connect time

2) Added flag OPT_NumbersAsStrings which, if set, returns all numbers as strings rather than integers or floats; this flag is used when defined variables are created (during select statements only)

3) Added flag OPT_Threading which, if set, uses OCI threading mode; there is a significant performance degradation in this mode (about 15-20%) but it does allow threads to share connections (threadsafety level 2 according to the Python Database API 2.0); note that in order to support this, Oracle 8i or higher is now required

4) Added Py_BEGIN_ALLOW_THREADS and Py_END_ALLOW_THREADS pairs where applicable to support threading during blocking OCI calls

5) Added global method attach() to cx_Oracle to support attaching to an existing database handle (as provided by PowerBuilder, for example)

6) Eliminated the cursor method fetchbinds() which was used for returning the list of bind variables after execution to get the values of out variables; the cursor method setinputsizes() was modified to return the list of bind variables and the cursor method execute() was modified to return the list of defined variables in the case of a select statement being executed; these variables have three methods available to them: getvalue([<pos>]) to get the value of a variable, setvalue(<pos>, <value>) to set its value and copy(<var>, <src_pos>, <targ_pos>) to copy the value from a variable in a more efficient manner than setvalue(getvalue())

7) Implemented cursor method executemany() which expects a list of dictionaries for the parameters

8) Implemented cursor method callproc()

9) Added cursor method prepare() which parses (prepares) the statement for execution; subsequent execute() or executemany() calls can pass None as the statement which will imply use of the previously prepared statement; used for high performance only

10) Added cursor method fetchraw() which will perform a raw fetch of the cursor returning the number of rows thus fetched; this is used to avoid the overhead of generating result sets; used for high performance only

11) Added cursor method executemanyprepared() which is identical to the method executemany() except that it takes a single parameter which is the number of times to execute a previously prepared statement and it assumes that the bind variables already have their values set; used for high performance only

12) Added support for rowid being returned in a select statement

13) Added support for comparing dates returned by cx_Oracle

14) Integrated patch from Andre Reitz to set the null ok flag in the description attribute of the cursor

15) Integrated patch from Andre Reitz to setup.py to support compilation with Python 1.5

16) Integrated patch from Benjamin Kearns to setup.py to support compilation on Cygwin

13.4.20 Version 2.4 (January 2002)

1) String variables can now be made any length (previously restricted to the 64K limit imposed by Oracle for default binding); use the type cx_Oracle.LONG_STRING as the parameter to setinputsizes() for binding in string values larger than 4000 bytes.

2) Raw and long raw columns are now supported; use the types cx_Oracle.BINARY and cx_Oracle.LONG_BINARY as the parameter to setinputsizes() for binding in values of these types.

3) Functions DateFromTicks(), TimeFromTicks() and TimestampFromTicks() are now implemented.

4) Function cursor.setoutputsize() implemented

5) Added the ability to bind arrays as out parameters to procedures; use the format [cx_Oracle.<DataType>, <NumElems>] as the input to the function setinputsizes() for binding arrays

6) Discovered from the Oracle 8.1.6 version of the documentation of the OCI libraries, that the size of the memory location required for the precision variable is larger than the printed documentation says; this was causing a problem with the code on the Sun platform.

7) Now support building RPMs for Linux.

13.4.21 Version 2.3 (October 2001)

1) Incremental performance enhancements (dealing with reusing cursors and bind handles)

2) Ensured that arrays of integers with a single float in them are all treated as floats, as suggested by Martin Koch.

3) Fixed code dealing with scale and precision for both defining a numeric variable and for providing the cursor description; this eliminates the problem of an underflow error (OCI-22054) when retrieving data with non-zero scale.

13.4.22 Version 2.2 (July 2001)

1) Upgraded thread safety to level 1 (according to the Python DB API 2.0) as an internal project required the ability to share the module between threads.

2) Added ability to bind ref cursors to PL/SQL blocks as requested by Brad Powell.

3) Added function write(Value, [Offset]) to LOB variables as requested by Matthias Kirst.

4) Procedure execute() on Cursor objects now permits a value None for the statement which means that the previously prepared statement will be executed and any input sizes set earlier will be retained. This was done to improve the performance of scripts that execute one statement many times.

5) Modified module global constants BINARY and DATETIME to point to the external representations of those types so that the expression type(var) == cx_Oracle.DATETIME will work as expected.

6) Added global constant version to provide means of determining the current version of the module.

7) Modified error checking routine to distinguish between an Oracle error and invalid handles.

8) Added error checking to avoid setting the value of a bind variable to a value that it cannot support and raised an exception to indicate this fact.

9) Added extra compile arguments for the AIX platform as suggested by Jehwan Ryu.

10) Added section to the README to indicate the method for a binary installation as suggested by Steve Holden.

11) Added simple usage example as requested by many people.

12) Added HISTORY file to the distribution.

License

Indices and tables

- genindex
- modindex
- search

Python Module Index

C

cx_Oracle, 13

Index

Symbols

H

I

K

L

M

N

O