

Qube 8.0 Complete Release Notes

```
#####
#
# Qube Release Notes
#
#####

#####
@RELEASE: 8.0-0

==== CL 23707 ====
@NEW: code to add a line that allows "qube" user to connect to DB from any machine on the network, with its
password

See also CL23705,23706

==== CL 23706 ====
@NEW: add qubedb_prep_0055.sql that sets the "qube" DB user's default password.

See also CL23705

==== CL 23705 ====
@CHANGE: Assign default password for "qube" DB user.
* Allow, in pg_hba.conf, access of "qube" user with password from non-localhost
* Make UI connect to DB over the network using the "qube" user

==== CL 23698 ====
@NEW: add python 3.9 and 3.10 support on macOS

==== CL 23680 ====
@FIX: TypeError when options.chunk unspecified

==== CL 23635 ====
@NEW: upgrading DB version to 55-- add qubedb_0055.sql .

==== CL 23608 ====
@CHANGE: make sure the *_sk columns (synthetic keys) in the fact tables are of "integer" data type instead
of "smallint"

==== CL 23600 ====
@NEW: add datawh/upgrade_scripts/upgrade_v4.sql

* also add datawh/alter_sk_to_int.sql which is included by the upgrade_v4.sql script to convert all *_sk
columns from smallint to int data type

ZD: 21636

==== CL 23596 ====
@NEW: add Shotgun (Shotgun) Python API 3.3.3 to qube-core package.

==== CL 23591 ====
@NEW: add Python API routine qb.jobtagvalues() that fetches and returns a list of existing values for job
tags (such as for "prod_show", "prod_seq", ..., "prod_custom1",...)

JIRA: QUBE-3924

==== CL 23589 ====
@FIX: datawh/upgrade_scripts/upgrade_v2.sql hadn't "SET ROLE dwh", resulting in some DBs created with
incorrect ownership

==== CL 23560 ====
@NEW: add C++ API routine "qb.jobtagvalues()" to fetch a list of existing values for job tags (such as for
"prod_show", "prod_seq", ..., "prod_custom1",...)

JIRA: QUBE-3924

==== CL 23516 ====
@NEW: Add Python 3.10 Qube API support (Linux and Windows)

JIRA: QUBE-3928

==== CL 23512 ====
@NEW: Add python 3.9 Qube API support (Linux and Windows)
```

JIRA: QUBE-3926

==== CL 23493 ====

@NEW: add API routines "qbbannedworkers()" (C++) and "qb.bannedworkers()" (Python) to return a list of all currently banned (removed/blacklisted) worker hosts.

* also added a "--listbanned" option, to "qbadmin w".

JIRA: QUBE-2074

==== CL 23475 ====

@CHANGE: Upping DATA_WAREHOUSE_VERSION to 3 in file "datawarehouse_version"

JIRA: QUBE-3920

==== CL 23473 ====

@NEW: add datawh/upgrade_scripts/upgrade_v3.sql . So far, adds calls to new external .sql scripts that create the new DWH subjob and work fact tables and initially populate them.

JIRA: QUBE-3920

==== CL 23472 ====

@NEW: dwh: add work fact table and related SQL code to populate it initially and on a regular basis.

JIRA: QUBE-3920

==== CL 23471 ====

@NEW: dwh: add subjob fact table and related SQL code to populate it initially and on a regular basis.

JIRA: QUBE-3920

==== CL 23470 ====

@FIX: add code to TRUNCATE the worker_dim and cluster_dim tables in populate_slotCount.sql where it initially populates those tables

==== CL 23453 ====

@FIX: tab char used for indentation in Python "helloWorld" jobtype example

==== CL 23419 ====

@CHANGE: add support for Perl v5.30 (especially for macOS 11.3.1+, and 11.5.2+) in api/perl/qb/qb.pm

ZD: 21313

JIRA: QUBE-3916

==== CL 23385 ====

@FIX: Message lengths over 254 resulting in Null characters being sent

==== CL 23363 ====

@CHANGE: removed Qt dependencies from qube core (especially the API)

JIRA: QUBE-3902

ZD: 21232, 21293

==== CL 23341 ====

@NEW: add (back) pyWorkCmdline jobtype, now python3 compatible

JIRA: QUBE-3896

ZD: 21197

==== CL 23340 ====

@CHANGE: convert python-based jobtype pyWorkCmdline from python2 to python3

==== CL 23324 ====

@FIX: data warehouse (dwh) upgrade_v2.sql script had issues with schema scoping, causing supervisor installation to prematurely fail

==== CL 23323 ====

@FIX: syntax error with the upgrade_v2.sql script for datawh, resulting in partially failed supervisor installation

==== CL 23290 ====

@CHANGE: support Apple-provided Python 3.8, instead of Home Brew version

==== CL 23262 ====

@CHANGE: postgresql startup scripts no longer control where the logs go (now it's specified in postgresql.conf)

JIRA: QUBE-3888

==== CL 23261 ====

@NEW: Add Qube-custom logging parameters to postgresql.conf. Now PGSQL logs are written to DATADIR/pg_log/pgsql.log on ALL platforms.

Formerly, log wasn't written to a file on Windows.

JIRA: QUBE-3888

==== CL 23260 ====

@FIX: some parsing issues in configure_postgresql_conf.py script

* bug where commented out parameters would end up in the end section of the postgresql.conf file, unless the commented out value exactly matches our value

* fixed regex to match better (had issues with params with empty value '', and with whitespace)

==== CL 23259 ====

@UPDATE: add signal handler for segfault in QbPreForkDaemon to print stack trace upon the said signal

==== CL 23258 ====

@FIX: a couple of issues with python scripts that export MySQL data and import it into PostgreSQL

* export script: quoting issue with the path to MySQL executable

* import script: jobid data must be imported before job data to satisfy a constraint added since 7.5-0

ZD: 21118

==== CL 23253 ====

@FIX: binarySort Python 3 compatibility

==== CL 23227 ====

@FIX: proxy program (proxy.exe) crashing under certain Windows environments

ZD: 21090

==== CL 23203 ====

@UPDATE: Edits for database_checks.py

==== CL 23201 ====

@FIX: Encoding xdrlib integers 128 and greater

==== CL 23163 ====

@CHANGE: make sure that the "pfx" db is created with its Encoding set to 'UTF8'

JIRA: QUBE-3865

==== CL 23160 ====

@NEW: Adding back the slotcount_fact table as well as the cached "data subset" tables to the Data Warehouse (dwh)

JIRA: QUBE-3867

ZD: 20996

* Add back, to dwh, the slotcount_fact table and all related .sql scripts:

create_hostState_dim.sql
create_slotCount_fact.sql
create_worker_dim.sql
populate_slotCount.sql
regular_slotCount.sql

* Also add upgrade_scripts/upgrade_v2.sql, and updated datawarehouse_version to 2.

In upgrade_v2.sql, reimplemented the PFX_CREATE_DATASUBSET_TABLE function (that was available in our MySQL dwh), and also added commands to source the create*.sql and populate_slotCount.sql files mentioned above to build the necessary tables.

* Add back cron jobs (Linux, macOS) and scheduledTask (Windows) that periodically run the regular_slotCount.sql collector, and build the "data subset" tables that contain data for a limited time-range (12hr, 36hr, 7day, 3wk, and 3mo) for faster preset queries for charting:

Linux: qube/etc/cron.d/com.pipelinefx.DataWarehouse.cron

macOS: add back qube/datawh/data_collectors/osx*.sh scripts in the data_collectors subdir, and their corresponding macOS cron drivers, qube/etc/com.pipelinefx.DataWarehouse.*.plist, qube/pkg/supepkg.pl and qu
bepkg.pm to package them up into the supervisor installer

Windows: changes made to [qubemsi.pm](#) and `rrd_tables.bat` to install and enable the scheduledTasks when the supervisor MSI installer is run

==== CL 23146 ====

@FIX: Python API: `qb.rangepartition()`, and consequently `qb.genpartitions()`, broken with python3, returning an empty list to valid input.

ZD: 20992

JIRA: QUBE-3863

==== CL 23052 ====

@FIX: include file path to the python execution module in the call to `exec()` so that `qb.backend.utils.getModulePath()` works properly

Also modified internals of `getModulePath()` in `qb/backend/utils.py` and `qb/util/__init__.py` to be more compatible across python versions

ZD: 20941

==== CL 23034 ====

@FIX: wildcard and regular expression (regex) filtering features in tools such as "qbjobs" broken since 7.0

The filtering features now work as described in doc at:

<http://docs.pipelinefx.com/display/QUBE/Using+Wildcards+and+Regular+Expressions>

JIRA: QUBE-3844

==== CL 23022 ====

@FIX: Python-based jobtype backends make undesirable calls to `qb.reportjob()`, rendering instance-based postflight return values useless

`pyCmdrange`, `pyCmdline`, `appFinder`, and `pyCmdrangeGPU` jobtypes were making a call to `qb.reportjob()` at the end of their execute.py scripts which is undesirable, and causing the return values of instance-based postflights to be meaningless. For example, a job instance could "complete", but the postflight that runs right after it could return non-zero indicating that the instance should be marked "failed". However, since the job instance was reporting "complete" via its own call to `qb.reportjob()` before running the postflight, the status that was amended by the postflight never took effect (it was too late to reach the supe).

JIRA: QUBE-3836

==== CL 22906 ====

@NEW:Add VRED 2021 support

==== CL 22873 ====

@FIX: additional fixes for "DataWH collectors aren't run by cron on Ubuntu"

JIRA: QUBE-3846

==== CL 22872 ====

@FIX: DataWH collectors aren't run by cron on Ubuntu

JIRA: QUBE-3846

==== CL 22845 ====

@FIX: add openssl 1.1.1h lib files to qube core package, which are needed by the Qt 5.14.2 network module used by the DRA.

JIRA: QUBE-3833

==== CL 22828 ====

@FIX: supervisor's embedded python3 interpreter fails running callbacks with error: `ModuleNotFoundError: No module named '_struct'`

ZD: 20850

JIRA: QUBE-3832

==== CL 22748 ====

@FIX: issue where some files/folders won't be deleted upon MSI uninstall.

==== CL 22735 ====

@FIX: Uninstalling supervisor on Windows now removes the postgresql software (but preserves the data directory).

@FIX: Windows: the 7zip self-extraction of PostgreSQL software throws an "ERROR: can not delete output file" error when upgrading supervisor

@CHANGE: Installation of supervisor on Windows now first makes sure that the previous installation of postgresql is removed.

JIRA: QUBE-3819, QUBE-3817

ZD: 20811, 20816

#####