

---

# MySQL Shell 8.2 Release Notes

## Abstract

This document contains release notes for the changes in MySQL Shell 8.2.

For additional MySQL Shell documentation, see <http://dev.mysql.com/>.

Updates to these notes occur as new product features are added, so that everybody can follow the development process. If a recent version is listed here that you cannot find on the download page (<https://dev.mysql.com/downloads/>), the version has not yet been released.

The documentation included in source and binary distributions may not be fully up to date with respect to release note entries because integration of the documentation occurs at release build time. For the most up-to-date release notes, please refer to the online documentation instead.

For legal information, see the [Legal Notices](#).

For help with using MySQL, please visit the [MySQL Forums](#), where you can discuss your issues with other MySQL users.

Document generated on: 2023-11-27 (revision: 27600)

## Table of Contents

Preface and Legal Notices .....	1
Changes in MySQL Shell 8.2.1 (2023-11-21, Innovation Release) .....	3
Changes in MySQL Shell 8.2.0 (2023-10-25, Innovation Release) .....	3

## Preface and Legal Notices

This document contains release notes for the changes in MySQL Shell 8.2.

### Legal Notices

Copyright © 1997, 2023, Oracle and/or its affiliates.

#### License Restrictions

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

#### Warranty Disclaimer

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

#### Restricted Rights Notice

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

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

### **Hazardous Applications Notice**

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

### **Trademark Notice**

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

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

### **Third-Party Content, Products, and Services Disclaimer**

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

### **Use of This Documentation**

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how Oracle disseminates it (that is, electronically for download on a Web site with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated

together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Oracle. Oracle and/or its affiliates reserve any and all rights to this documentation not expressly granted above.

## Documentation Accessibility

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

## Access to Oracle Support for Accessibility

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

## Changes in MySQL Shell 8.2.1 (2023-11-21, Innovation Release)

### Utilities Bugs Fixed

- It was not possible to migrate from a compatible database using the `copyInstance()` or `dumpInstance()` utilities. Errors similar to the following error were returned:

```
Util.copyInstance: Unknown system variable 'activate_all_roles_on_login' (MYSQLSH 1193)
```

(Bug #35963431)

## Changes in MySQL Shell 8.2.0 (2023-10-25, Innovation Release)

- [Deprecation and Removal Notes](#)
- [AdminAPI Added or Changed Functionality](#)
- [AdminAPI Bugs Fixed](#)
- [Utilities Added or Changed Functionality](#)
- [Utilities Bugs Fixed](#)
- [Bugs Fixed](#)

## Deprecation and Removal Notes

- The following are deprecated in this version and subject to removal in a future version of MySQL Shell:
  - `ociparmanifest`: this dump utility option has been superceded by more efficient methods of creating and using Pre-Authenticated Requests.
  - `--import`: superceded by the `importJson` utility, which is also available from the command line.
  - `--recreate-schema`: no longer required.

Deprecation warnings will be displayed when any of these are used. (WL #15863)

- AdminAPI support for MySQL Server 5.7 is deprecated and subject to removal in a future version. AdminAPI commands will emit a deprecation warning if run against a 5.7 instance.

The following AdminAPI methods are deprecated in this release, and subject to removal in a future version:

- `dba.configureLocalInstance()`: this method is only required by MySQL versions which did not support remote persisting of variables. MySQL Server 5.7 is currently the only supported version with that limitation.
- `Cluster.checkInstanceState()`: the checks implemented by this method are now integrated into methods which require it, such as `Cluster.addInstance()`, `Cluster.rejoinInstance()`, and so on.

The following AdminAPI options are deprecated and subject to removal in a future version:

- `interactive` is deprecated on all AdminAPI methods which use it.
- `password` is deprecated on the following methods:
  - `dba.rebootClusterFromCompleteOutage()`
  - `dba.configureInstance()`
  - `dba.configureLocalInstance()`
  - `dba.configureReplicaSetInstance()`
  - `Cluster.checkInstanceConfiguration()`
  - `Cluster.removeInstance()`
  - `Cluster.addInstance()`
  - `Cluster.forceQuorumUsingPartitionOf()`
  - `Cluster.rejoinInstance()`
- `waitRecovery` is deprecated on all AdminAPI methods which use it, use `recoveryProgress` instead.

The `recoveryProgress` option is also added to `Cluster.addInstance()`, `ReplicaSet.addInstance()`, and `ReplicaSet.rejoinInstance()`.

- `clearReadOnly` option is deprecated in `dba.configureInstance()`, `dba.configureLocalInstance()`, and `dba.dropMetadataSchema`.

`super_read_only` is checked, and automatically disabled.

Deprecation warnings will be emitted in MySQL Shell 8.2 for each of these deprecated features, methods, and options. (WL #15862)

## AdminAPI Added or Changed Functionality

- AdminAPI now supports the configuration of the following replication options:
  - `SOURCE_CONNECT_RETRY` can be configured using the following options:

- `clusterSetReplicationConnectRetry` of `clusterSet.createReplicaCluster()` and `cluster.setOption()`.
- `replicationConnectRetry` of `replicaSet.addInstance()` and `replicaSet.setInstanceOption()`.
- `SOURCE_RETRY_COUNT` can be configured using the following options:
  - `clusterSetReplicationRetryCount` of `clusterSet.createReplicaCluster()` and `cluster.setOption()`.
  - `replicationRetryCount` of `replicaSet.addInstance()` and `replicaSet.setInstanceOption()`.
- `SOURCE_HEARTBEAT_PERIOD` can be configured using the following options:
  - `clusterSetReplicationHeartbeatPeriod` of `clusterSet.createReplicaCluster()` and `cluster.setOption()`.
  - `replicationHeartbeatPeriod` of `replicaSet.addInstance()` and `replicaSet.setInstanceOption()`.
- `SOURCE_COMPRESSION_ALGORITHMS` can be configured using the following options:
  - `clusterSetReplicationCompressionAlgorithms` of `clusterSet.createReplicaCluster()` and `cluster.setOption()`.
  - `replicationCompressionAlgorithms` of `replicaSet.addInstance()` and `replicaSet.setInstanceOption()`.
- `SOURCE_ZSTD_COMPRESSION_LEVEL` can be configured using the following options:
  - `clusterSetReplicationZstdCompressionLevel` of `clusterSet.createReplicaCluster()` and `cluster.setOption()`.
  - `replicationZstdCompressionLevel` of `replicaSet.addInstance()` and `replicaSet.setInstanceOption()`.
- `SOURCE_BIND` can be configured using the following options:
  - `clusterSetReplicationBind` of `clusterSet.createReplicaCluster()` and `cluster.setOption()`.
  - `replicationBind` of `replicaSet.addInstance()` and `replicaSet.setInstanceOption()`.
- `NETWORK_NAMESPACE` can be configured using the following options:
  - `clusterSetReplicationNetworkNamespace` of `clusterSet.createReplicaCluster()` and `cluster.setOption()`.
  - `replicationNetworkNamespace` of `replicaSet.addInstance()` and `replicaSet.setInstanceOption()`.

For more information on the replication options and their values, see [CHANGE REPLICATION SOURCE TO Statement](#).

`cluster.options` and `replicaSet.options` are extended to display these options and their values for each instance in the cluster.

`clusterset.status` is extended for Replica Clusters to display an error in the `clusterErrors` array for any replication channel which is not correctly configured according to the Cluster's metadata. Similarly for `cluster.status` if the Cluster is a Replica Cluster in a ClusterSet. (WL #13687)

- The following methods are extended to include information on the MySQL Router Read/Write Splitting Port, `rwSplitPort`:

- `cluster.listRouters()`
- `clusterSet.listRouters()`
- `replicaSet.listRouters()`

(WL #15552)

- `cluster.setRoutingOption()` is extended with the option `unreachable_quorum_allowed_traffic`. This option enables configuration of MySQL Router's routing policy in the event of a loss of quorum on the only reachable Cluster partition.

See [Routing Options](#).

`cluster.routingOptions()` is extended to list the value of `unreachable_quorum_allowed_traffic`. (WL #15842)

## AdminAPI Bugs Fixed

- Using `-- clusterset listRouters` on the command line, without providing a parameter for `listRouters`, resulted in the following error in MySQL Shell 8.0.34:

```
ERROR: Argument #1 is expected to be a string
```

(Bug #35747208)

References: This issue is a regression of: Bug #35068427.

- If multiple Read Replicas are added with the same label, only one of the Read Replicas was listed in the `Cluster.status()` output.

As of this release, Read Replica labels must be unique within the Cluster and can only contain alphanumeric values, or `_` (underscore), `.` (period), `-` (hyphen), or `:` (colon) characters. If the label is not unique within the Cluster, or invalid characters are detected, an error is returned. (Bug #35739776)

- If `memberAuthType` was set to `CERT_ISSUER` or `CERT_SUBJECT`, the values defined for the ClusterSet option `clusterSetReplicationSslMode` and ReplicaSet option `replicationSslMode` were ignored and automatically set to `VERIFY_CA`. (Bug #35621465)
- `Cluster.setPrimaryInstance()` did not check if the target instance was the primary of a Replica Cluster, resulting in a misconfigured replication channel. An error similar to the following was returned:

```
WARNING: The Cluster's Replication Channel is misconfigured or stopped,
topology changes will not be allowed on the InnoDB Cluster
'clusterName'
NOTE: To restore the Cluster and ClusterSet operations, the Replication
Channel must be fixed using rejoinCluster()
```

As of this release, `Cluster.setPrimaryInstance()` checks if the target instance is a configured primary and stops if it is, returning an error to the user. (Bug #35594376)

- If a switchover occurred in a ClusterSet where one of the replica clusters was entirely offline, attempting to reboot the cluster produced errors relating to rejoining the offline cluster to the ClusterSet.

As of this release, no attempt is made to rejoin the offline cluster to the ClusterSet and warnings are displayed to advise the user that the replica cluster is invalid and must be rejoined manually. (Bug #35548072)

- `replicationSslMode` was not returned by `ReplicaSet.status()` or `ReplicaSet.options()`. (Bug #35523549)
- If a Replica Cluster with Read Replicas had a major outage, `ClusterSet.status({extended:1})` showed the Read Replicas as Cluster members with the `memberRole` of `SECONDARY`. (Bug #35449470)
- If a Cluster was rebooted using the `dba.rebootClusterFromCompleteOutage()` method, without specifying the `switchCommunicationStack` option, the Cluster was rebooted using the defaults for the `xcom` communication stack, even if the Cluster was configured to use `mysql`. As a result, recovery accounts were not recreated and `mysql`-specific checks were not performed.

As of this release, `dba.rebootClusterFromCompleteOutage()` checks the configured communication stack and proceeds using the configured value. (Bug #35444206)

- `ReplicaSet.forcePrimaryInstance` could not use an instance with the applier state OFF as a potential primary. An error similar to the following was returned:

```
ERROR: Replication applier is OFF at instance instanceName:3306.
ERROR: Replication applier is OFF at instance instanceName:3306.
ERROR: Replication errors found for one or more SECONDARY instances. Use the
'invalidateErrorInstances' option to perform the failover anyway by skipping
and invalidating instances with errors.
ReplicaSet.forcePrimaryInstance: One or more instances have replication
applier errors. (MYSQLSH 51145)
```

It was not possible to use such an instance even if `invalidateErrorInstances` was set to `true`.

As of this release, `ReplicaSet.forcePrimaryInstance` treats such instances as valid candidates for promotion. `invalidateErrorInstances` is not required and the command does not wait for transactions to be applied to the instance. (Bug #35434959)

- It was not possible to retrieve a Cluster handle, using `dba.getCluster()`, from a Read Replica connected to a Cluster which had lost quorum. An error similar to the following was returned:

```
Dba.getCluster: Dba.getCluster: Unable to find a cluster PRIMARY member from
the active shell session because the cluster has too many UNREACHABLE members
and no quorum is possible.
Use Dba.getCluster(null, {connectToPrimary:false}) to get a read-only cluster
handle. (RuntimeError)
```

As of this release, `dba.getCluster()` establishes a connection to a Cluster member and returns the Cluster handle. (Bug #35391093)

- It is not possible for an instance to join a Cluster if the instance's local address, as generated by the AdminAPI or defined in `localAddress`, is not compatible with the automatically-generated Group Replication allow list.

As of this release, `dba.createCluster()` and `cluster.addInstance()` verify the local address (supplied or generated) is compatible with the automatically generated allow list. If the list is manually defined, no test is performed.

**Note**

This only applies if the communication stack is XCOM.

Additionally, on Windows platforms, if the allow list is not automatically generated, a warning is issued regarding the server trying to connect to itself. (Bug #31357039)

## Utilities Added or Changed Functionality

- Instance dump utility now excludes the `mysql_firewall` schema if `ocimds: true`. (Bug #35805866)
- MySQL Shell was updated for compatibility with the privilege changes made in MySQL HeatWave Service.

The following privileges were added to MySQL HeatWave Service:

- `AUDIT_ADMIN`
- `BACKUP_ADMIN`
- `FLUSH_OPTIMIZER_COSTS`
- `FLUSH_STATUS`
- `FLUSH_TABLES`
- `FLUSH_USER_RESOURCES`
- `ROLE_ADMIN`

The following privileges were removed from MySQL HeatWave Service

- `RESOURCE_GROUP_ADMIN`
- `RESOURCE_GROUP_USER`

For more information on MySQL HeatWave Service privileges, see [Default MySQL Privileges](#). (Bug #35668544)

- Operations resulting in the curl errors `52: CURLE_GOT_NOTHING` and `56: CURLE_RECV_ERROR` are retried for all supported cloud vendors, for all utilities which support them. (Bug #35659057)
- MySQL Shell utilities which connect to Object Storage now use dedicated endpoints.

For more information, see [OCI Object Storage Dedicated Endpoints](#). (Bug #35561100)

- `util.checkForServerUpgrade()` has been updated to check for columns which have foreign keys referencing columns in tables using different database storage engines, such as MyISAM. (Bug #35155064)
- The `@.done.json` generated by `util.dumpSchemas()` now includes the number of rows dumped for each table, in the `tableRows` field. (Bug #34904657)



- As of MySQL Server 8.2.0, `SET_USER_ID` is deprecated and subject to removal in a future version. `SET_USER_ID` is replaced by `SET_ANY_DEFINER` and `ALLOW_NONEXISTENT_DEFINER`. This change impacts the way MySQL Shell handles dumps for use in MySQL HeatWave Service (`ocimds: true`) because the administrator user will have the `SET_ANY_DEFINER` privilege and is now able to execute statements with the `DEFINER` clause. This was not possible in previous versions and users had to use the `strip_definers` compatibility option.

As of this release, a new option, `targetVersion` is added to the `util.dumpInstance()`, `util.dumpSchemas()`, and `util.dumpTables()` methods. You must use this method to manually define the version of the target MySQL instance, in n.n.n format. Such as 8.1.0, for example. If the value is not set, the MySQL Shell version is used.

The compatibility checks are adjusted depending on the value of `targetVersion`.



#### Note

This option is not required by the `util.copyInstance()`, `util.copySchemas()`, and `util.copyTables()` methods, which can detect the version of the target instance automatically.

If `ocids:true` and `targetVersion` is set to a version which supports the `SET_ANY_DEFINER`, the compatibility option `strip_definers` is not required.

If the compatibility option `strip_restricted_grants` is used, and `targetVersion` is set to a version which supports the `SET_ANY_DEFINER`, `SET_USER_ID` is replaced by `SET_ANY_DEFINER` instead of being stripped. (WL #15887)

## Utilities Bugs Fixed

- When dumping an instance, with `ocimds: true`, MySQL Shell reported warnings relating to users with privileges on system schemas. As these schemas are always present in both source and target, the warnings were unnecessary and have been removed. (Bug #35680824)
- Under certain circumstances, when loading a file larger than `maxBytesPerTransaction`, (or `1.5 * bytesPerChunk` if `maxBytesPerTransaction` was not used) a memory leak could occur. (Bug #35600174)
- When loading a very large file with `util.importTable()`, the progress was displayed as 100% complete although the process was not yet complete and the server continued to process the imported data. This could lead a user to believe the process had failed and result in manual attempts to cancel the process.

As of this release, thread activity is now displayed for the commit process, and an animated hyphen (-) is used to indicate progress. (Bug #35590038)

- Importing a large file to a non-existent schema, using `util.importTable()`, took a long time to return an unknown database error. This was caused by MySQL Shell chunking the entire file into smaller chunks without checking for errors until the entire file was chunked.

As of this release, errors generated by other threads are checked during the chunking process. (Bug #35541522)

- `util.debug.collectDiagnostics()` failed if run on an InnoDB Cluster created with an older version of MySQL Shell. An error similar to the following was generated:

```
An error occurred during data collection. Partial output deleted.
debug.collectDiagnostics: ClassicSession.run_sql: Table
```

```
'mysql_innodb_cluster_metadata.v2_cs_clustersets' doesn't exist (MySQL Error 1146)
```

As of this release, `util.debug.collectDiagnostics()` collects diagnostics information even if tables are missing. (Bug #35468106)

- Exporting to a bucket using a prefix with `util.exportTable()`, required the prefix to exist in the target bucket. The following error was returned if the prefix did not exist:

```
ValueError: Util.export_table: Cannot proceed with the dump, the directory containing 'prefix/filename' does not exist at the target location table-export.
```

As of this release, this requirement has been removed. If the prefix does not exist, it is created. (Bug #35462985)

- `util.debug.collectDiagnostics()` threw an exception if the server was configured to write error logs to stderr. For example, if `mysqld` was started with `--console`. (Bug #35318770)
- Under certain circumstances, `util.importTable()` failed when importing a single, uncompressed file due to how the chunking algorithm handled unescaped sequences of characters used by `fieldsTerminatedBy`.

As of this release, if a single file is imported, additional threads read the file in fixed-size blocks, then the blocks are scanned sequentially for row boundaries. Once read, the file data is stored in memory until loaded by the server. This also applies to compressed files which are now also chunked.

Additionally, chunking and sub-chunking is disabled if `linesTerminatedBy` is set to an empty string or the same value as `fieldsTerminatedBy`. (Bug #35279351)

- The MySQL Shell upgrade checker utility flagged views as corrupt in MySQL 5.7 versions up to 5.7.39. This issue occurred for views whose `from` clause contained a table schema prefix with a `group by` clause. (Bug #111813, Bug #35635009)
- `util.dumpInstance()` generated an empty `.pre.sql` file for invalid views.

As of this release, invalid views are detected and generate an error. Such views must be corrected or excluded from the dump using the `excludeTables` option. (Bug #111092, Bug #35415976)

- The utility, `checkForServerUpgrade`, did not recognize `INTERSECT` as a reserved word.

`INTERSECT` was reserved in MySQL Server 8.0.31. (Bug #110824, Bug #35335813)

## Bugs Fixed

- It was not possible to enter a full-width 8 (unicode 0xFF18) using copy and paste or a Japanese input method. (Bug #35782407)
- Running redirected commands similar to

```
mysqlsh.exe --pym pip install --upgrade pip > out.log
```

resulted in an exception containing the following:

```
File
"C:\path\mysql-shell\lib\Python3.10\lib\site-pack
ages\pip\_vendor\rich\console.py", line 1990, in _check_buffer
```

```
self.file.fileno() in _STD_STREAMS_OUTPUT
AttributeError: module 'mysqlsh.shell_stdout' has no attribute 'fileno'
```

As of this release, a `fileno()` method is exposed in MySQL Shell's `stdout` and `stderr` where it can also be found by redirected commands. (Bug #35528045)

- If an empty string was passed on the command line, MySQL Shell returned `stoi` or `invalid stoi argument`, depending on platform. For example:

```
>mysqlsh root@localhost?connect-timeout=
MySQL Shell 8.1.1-commercial
.....
Creating a Classic session to 'root@localhost?connect-timeout='
invalid stoi argument
```

As of this release, if parameters such as `connect-timeout` are provided without values on the command line, an error similar to the following is returned:

```
Invalid URI: Invalid value '' for 'connect-timeout'. The timeout value must be a positive integer (includ
```

(Bug #35130320)

