
Employees Sample Database

Table of Contents

1 Preface and Legal Notices	1
2 Introduction	2
3 Installation	3
4 Validating the Employee Data	4
5 Employees Structure	5
6 License for the Employees Database	7

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: 2021-08-05 (revision: 70529)

1 Preface and Legal Notices

This document describes the `employees` sample database—its history, installation, structure and usage.

Legal Notices

Copyright © 2008, 2021, Oracle and/or its affiliates.

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

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

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

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

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

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

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

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

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.

Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

2 Introduction

This document describes the Employees sample database.

The Employees sample database was developed by Patrick Crews and Giuseppe Maxia and provides a combination of a large base of data (approximately 160MB) spread over six separate tables and consisting of 4 million records in total. The structure is compatible with a wide range of storage engine types. Through an included data file, support for partitioned tables is also provided.

In addition to the base data, the Employees database also includes a suite of tests that can be executed across the test data to ensure the integrity of the data that you have loaded. This should help ensure the quality of the data during initial load, and can be used after usage to ensure that no changes have been made to the database during testing.

3 Installation

The Employees database is available from [Employees DB on GitHub](#). You can download a prepackaged archive of the data, or access the information through Git.

To use the Zip archive package, download the archive and unpack it using [WinZip](#) or another tool that can read `.zip` files, then change location into the unpacked package directory. For example, using `unzip`, execute these commands:

```
shell> unzip test_db-master.zip
shell> cd test_db-master/
```

The Employees database is compatible with several different storage engines, with the InnoDB engine enabled by default. Edit the `employees.sql` file and adjust the comments to choose a different storage engine:

```
set storage_engine = InnoDB;
-- set storage_engine = MyISAM;
-- set storage_engine = Falcon;
-- set storage_engine = PBXT;
-- set storage_engine = Maria;
```

To import the data into your MySQL instance, load the data through the `mysql` command-line tool:

```
shell> mysql -t < employees.sql
+-----+
| INFO |
+-----+
| CREATING DATABASE STRUCTURE |
+-----+
+-----+
| INFO |
+-----+
| storage engine: InnoDB |
+-----+
+-----+
| INFO |
+-----+
| LOADING departments |
+-----+
+-----+
| INFO |
+-----+
| LOADING employees |
+-----+
+-----+
| INFO |
+-----+
| LOADING dept_emp |
+-----+
+-----+
| INFO |
+-----+
| LOADING dept_manager |
+-----+
+-----+
| INFO |
+-----+
| LOADING titles |
+-----+
+-----+
| INFO |
+-----+
| LOADING salaries |
```

To test that the data you have loaded matches the expected results, run the test suite. For more information, see [Section 4, “Validating the Employee Data”](#).

4 Validating the Employee Data

You can validate the Employee data using two methods, `md5` and `sha`. Two SQL scripts are provided for this purpose, `test_employees_sha.sql` and `test_employees_md5.sql`. To run the tests, use `mysql`:

```
shell>time mysql -t < test_employees_sha.sql
+-----+
| INFO |
+-----+
| TESTING INSTALLATION |
+-----+
+-----+
| table_name | expected_records | expected_crc |
+-----+
| employees | 300024 | 4d4aa689914d8fd41db7e45c2168e7dcb9697359 |
| departments | 9 | 4b315afa0e35ca6649df897b958345bcb3d2b764 |
| dept_manager | 24 | 9687a7d6f93ca8847388a42a6d8d93982a841c6c |
| dept_emp | 331603 | f16f6ce609d032d6b1b34748421e9195c5083da8 |
| titles | 443308 | d12d5f746b88f07e69b9e36675b6067abb01b60e |
| salaries | 2844047 | b5a1785c27d75e33a4173aaa22ccf41ebd7d4a9f |
+-----+
+-----+
| table_name | found_records | found_crc |
+-----+
| employees | 300024 | 4d4aa689914d8fd41db7e45c2168e7dcb9697359 |
| departments | 9 | 4b315afa0e35ca6649df897b958345bcb3d2b764 |
| dept_manager | 24 | 9687a7d6f93ca8847388a42a6d8d93982a841c6c |
| dept_emp | 331603 | f16f6ce609d032d6b1b34748421e9195c5083da8 |
| titles | 443308 | d12d5f746b88f07e69b9e36675b6067abb01b60e |
| salaries | 2844047 | b5a1785c27d75e33a4173aaa22ccf41ebd7d4a9f |
+-----+
+-----+
| table_name | records_match | crc_match |
+-----+
| employees | OK | ok |
| departments | OK | ok |
| dept_manager | OK | ok |
| dept_emp | OK | ok |
| titles | OK | ok |
| salaries | OK | ok |
+-----+

real 0m37.067s
user 0m0.007s
sys 0m0.009s

shell>time mysql -t < test_employees_md5.sql
+-----+
| INFO |
+-----+
| TESTING INSTALLATION |
+-----+
+-----+
| table_name | expected_records | expected_crc |
+-----+
| employees | 300024 | 4ec56ab5ba37218d187cf6ab09ce1aa1 |
| departments | 9 | dlaf5e170d2d1591d776d5638d71fc5f |
| dept_manager | 24 | 8720e2f0853ac9096b689c14664f847e |
| dept_emp | 331603 | c2c4fc7f0506e50959a6c67ad55cac31 |
+-----+
```

```


| titles      |          443308 | bfa016c472df68e70a03facafalbc0a8 |
| salaries   |          2844047 | fd220654e95aea1b169624ffe3fca934 |
+-----+-----+-----+
| table_name | found_records  | found_crc                            |
+-----+-----+-----+
| employees  |          300024 | 4ec56ab5ba37218d187cf6ab09celaal |
| departments|              9 | dlaf5e170d2d1591d776d5638d71fc5f |
| dept_manager|             24 | 8720e2f0853ac9096b689c14664f847e |
| dept_emp   |          331603 | c2c4fc7f0506e50959a6c67ad55cac31 |
| titles     |          443308 | bfa016c472df68e70a03facafalbc0a8 |
| salaries   |          2844047 | fd220654e95aea1b169624ffe3fca934 |
+-----+-----+-----+
| table_name | records_match | crc_match |
+-----+-----+-----+
| employees  | OK            | ok        |
| departments| OK            | ok        |
| dept_manager| OK            | ok        |
| dept_emp   | OK            | ok        |
| titles     | OK            | ok        |
| salaries   | OK            | ok        |
+-----+-----+-----+

real 0m33.453s
user 0m0.007s
sys 0m0.009s

```

5 Employees Structure

The following diagram provides an overview of the structure of the Employees sample database.



6 License for the Employees Database

This work is licensed under the Creative Commons Attribution-Share Alike 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

Please contact <http://www.mysql.com/about/contact/> for more information.

