

MySQL 8.0.12 InnoDB 表空间可视化工具 innodb_ruby

此工具为第三方开源的表空间可视化工具，详见文末的参考文献链接。
下述的演示安装环境是 CentOS Linux 7

1. 安装软件 ruby

```
yum -y install ruby rubygems
```

2. 更改下载源为 ruby-china

```
gpg2 --keyserver hkp://pool.sks-keyservers.net --recv-keys  
409B6B1796C275462A1703113804BB82D39DC0E3  
7D2BAF1CF37B13E2069D6956105BD0E739499BDB  
gem sources -a http://mirrors.aliyun.com/rubygems/  
or  
$ gem sources --add https://gems.ruby-china.com/ --remove https://rubygems.org/  
$ gem sources -l  
https://gems.ruby-china.com  
curl -sSL https://get.rvm.io | bash -s stable  
source /etc/profile.d/rvm.sh  
rvm list known  
rvm install 2.7  
ruby -v
```

3. 安装 innodb_ruby

```
# gem install innodb_ruby  
Fetching: bindata-2.4.3.gem (100%)  
Successfully installed bindata-2.4.3  
Fetching: digest-crc-0.4.1.gem (100%)  
Successfully installed digest-crc-0.4.1  
Fetching: innodb_ruby-0.9.15.gem (100%)  
Successfully installed innodb_ruby-0.9.15  
Parsing documentation for bindata-2.4.3  
Installing ri documentation for bindata-2.4.3  
Parsing documentation for digest-crc-0.4.1  
Installing ri documentation for digest-crc-0.4.1  
Parsing documentation for innodb_ruby-0.9.15  
Installing ri documentation for innodb_ruby-0.9.15  
gems installed
```

4. innodb_ruby 使用命令简介

```
[root@niegq mysql]# innodb_space --help  
Usage: innodb_space <options> <mode>
```

常用参数示例:

`innodb_space -s ibdata1 [-T table-name [-I index-name [-R record-offset]]] [options]`

`innodb_space -f file-name.ibd [-r ./descriptor.rb -d DescriptorClass] [options] <mode>`

Use the file-name.ibd tablespace file (and the DescriptorClass describer where required) to read the tablespace structures or indexes.

比如:

`innodb_space -f dbcourse/nieprofiles.ibd space-summary`

`innodb_space -f dbcourse/nieprofiles.ibd -F 3 space-index-fseg-pages-summary`

`innodb_space -f dbcourse/finalanswers_his.ibd space-inodes-summary`

部分参数选项:

`--help, -?`

Print this usage text.

`--trace, -t`

Enable tracing of all data read. Specify twice to enable even more tracing (including reads during opening of the tablespace) which can be quite noisy.

`--system-space-file, -s <arg>`

Load the system tablespace file or files <arg>: Either a single file e.g. 'ibdata1', a comma-delimited list of files e.g. 'ibdata1,ibdata1', or a directory name. If a directory name is provided, it will be scanned for all files named 'ibdata?' which will then be sorted alphabetically and used to load the system tablespace.

If using the `--system-space-file` option, the following options may also be used:

`--table-name, -T <name>`

Use the table name <name>.

`--index-name, -I <name>`

Use the index name <name>.

`--system-space-tables, -x`

Allow opening tables from the system space to support system spaces with tables created without `innodb-file-per-table` enabled.

`--data-directory, -D <directory>`

Open per-table tablespace files from <directory> rather than from the directory where the system-space-file is located.

`--space-file, -f <file>`

Load the tablespace file <file>.

- page, -p <page>
Operate on the page <page>.

- record, -R <offset>
Operate on the record located at <offset> within the index page.

- level, -l <level>
Operate on the level <level>.

- list, -L <list>
Operate on the list <list>.

- fseg-id, -F <fseg_id>
Operate on the file segment (fseg) <fseg_id>.

- require, -r <file>
Use Ruby's 'require' to load the file <file>. This is useful for loading classes with record describers.

- describer, -d <describer>
Use the named record describer to parse records in index pages.

The following modes are supported:

- system-spaces
Print a summary of all spaces in the system.

- data-dictionary-tables
Print all records in the SYS_TABLES data dictionary table.

- data-dictionary-columns
Print all records in the SYS_COLUMNS data dictionary table.

- data-dictionary-indexes
Print all records in the SYS_INDEXES data dictionary table.

- data-dictionary-fields
Print all records in the SYS_FIELDS data dictionary table.

- space-summary
Summarize all pages within a tablespace. A starting page number can be provided with the --page/-p argument.

- space-index-pages-summary

Summarize all 'INDEX' pages within a tablespace. This is useful to analyze page fill rates and record counts per page. In addition to 'INDEX' pages, 'ALLOCATED' pages are also printed and assumed to be completely empty. A starting page number can be provided with the --page/-p argument.

space-index-fseg-pages-summary

The same as space-index-pages-summary but only iterate one fseg, provided with the --fseg-id/-F argument.

space-index-pages-free-plot

Use Ruby's gnuplot module to produce a scatterplot of page free space for all 'INDEX' and 'ALLOCATED' pages in a tablespace. More aesthetically pleasing plots can be produced with space-index-pages-summary output, but this is a quick and easy way to produce a passable plot. A starting page number can be provided with the --page/-p argument.

space-page-type-regions

Summarize all contiguous regions of the same page type. This is useful to provide an overall view of the space and allocations within it. A starting page number can be provided with the --page/-p argument.

space-page-type-summary

Summarize all pages by type. A starting page number can be provided with the --page/-p argument.

space-indexes

Summarize all indexes (actually each segment of the indexes) to show the number of pages used and allocated, and the segment fill factor.

space-lists

Print a summary of all lists in a space.

space-list-iterate

Iterate through the contents of a space list.

space-extents

Iterate through all extents, printing the extent descriptor bitmap.

space-extents-illustrate

Iterate through all extents, illustrating the extent usage using ANSI color and Unicode box drawing characters to show page usage throughout the space.

space-lsn-age-illustrate

Iterate through all pages, producing a heat map colored by the page LSN using ANSI color and Unicode box drawing characters, allowing the user to get an overview of page modification recency.

space-inodes-fseg-id

Iterate through all inodes, printing only the FSEG ID.

space-inodes-summary

Iterate through all inodes, printing a short summary of each FSEG.

space-inodes-detail

Iterate through all inodes, printing a detailed report of each FSEG.

index-recurse

Recurse an index, starting at the root (which must be provided in the first --page/-p argument), printing the node pages, node pointers (links), leaf pages. A record describer must be provided with the --describer/-d argument to recurse indexes (in order to parse node pages).

index-record-offsets

Recurse an index as index-recurse does, but print the offsets of each record within the page.

index-digraph

Recurse an index as index-recurse does, but print a dot-compatible digraph instead of a human-readable summary.

index-level-summary

Print a summary of all pages at a given level (provided with the --level/-l argument) in an index.

index-fseg-internal-lists

index-fseg-leaf-lists

Print a summary of all lists in an index file segment. Index root page must be provided with --page/-p.

index-fseg-internal-list-iterate

index-fseg-leaf-list-iterate

Iterate the file segment list (whose name is provided in the first --list/-L argument) for internal or leaf pages for a given index (whose root page is provided in the first --page/-p argument). The lists used for each index are 'full', 'not_full', and 'free'.

index-fseg-internal-frag-pages

index-fseg-leaf-frag-pages

Print a summary of all fragment pages in an index file segment. Index root page must be provided with --page/-p.

page-dump

Dump the contents of a page, using the Ruby pp ('pretty-print') module.

page-account

Account for a page's usage in FSEGs.

page-validate

Validate the contents of a page.

page-directory-summary

Summarize the record contents of the page directory in a page. If a record describer is available, the key of each record will be printed.

page-records

Summarize all records within a page.

page-illustrate

Produce an illustration of the contents of a page.

record-dump

Dump a detailed description of a record and the data it contains. A record offset must be provided with -R/--record.

record-history

Summarize the history (undo logs) for a record. A record offset must be provided with -R/--record.

undo-history-summary

Summarize all records in the history list (undo logs).

undo-record-dump

Dump a detailed description of an undo record and the data it contains. A record offset must be provided with -R/--record.

参考文献:

Refer to

https://github.com/jeremycollection/innodb_ruby/wiki

<https://blog.csdn.net/vkingnew/article/details/82775944>