


Creating INFORMATION_SCHEMA Tables

MySQL Conference & Expo 2007
Mike Kruckenberg
<http://mike.kruckenberg.com>

about me

(who I am)

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007



Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

about me

(who I am not)

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

this presentation

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

might be obsolete
(with pluggable I_S around the corner)

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

not totally irrelevant
(understaning native I_S is helpful)

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

what is I_S?

- set of tables that provide information about server and database schemas
- virtual, in memory, and read only
- made available in the `information_schema` database
- like `show`, but more
- standards-based system catalog
- query-based interface
- added in 5.0.2

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

I_S Tables

- CHARACTER_SETS
- COLLATIONS
- COLLATION_CHARACTER_SET_APPLICABILITY
- COLUMNS
- COLUMN_PRIVILEGES
- ENGINES
- EVENTS
- FILES
- GLOBAL_STATUS
- GLOBAL_VARIABLES
- KEY_COLUMN_USAGE
- PARTITIONS
- PLUGINS
- PROCESSLIST
- REFERENTIAL_CONSTRAINTS
- ROUTINES
- SCHEMATA
- SCHEMA_PRIVILEGES
- SESSION_STATUS
- SESSION_VARIABLES
- STATISTICS
- TABLES
- TABLE_CONSTRAINTS
- TABLE_PRIVILEGES
- TRIGGERS
- USER_PRIVILEGES
- VIEWS

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

I_S statistics Table

```
mysql> desc statistics;
```

Field	Type	Null	Key	Default	Extra
TABLE_CATALOG	varchar(512)	YES		NULL	
TABLE_SCHEMA	varchar(64)	NO			
TABLE_NAME	varchar(64)	NO			
NON_UNIQUE	bigint(1)	NO		0	
INDEX_SCHEMA	varchar(64)	NO			
INDEX_NAME	varchar(64)	NO			
SEQ_IN_INDEX	bigint(2)	NO		0	
COLUMN_NAME	varchar(64)	NO			
COLLATION	varchar(1)	YES		NULL	
CARDINALITY	bigint(21)	YES		NULL	
SUB_PART	bigint(3)	YES		NULL	
PACKED	varchar(10)	YES		NULL	
NULLABLE	varchar(3)	NO			
INDEX_TYPE	varchar(16)	NO			
COMMENT	varchar(16)	YES		NULL	

```
15 rows in set (0.01 sec)
```

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

Quick I_S Example

```
mysql> show tables;  
+-----+  
| Tables_in_sakila |  
+-----+  
| actor              |  
| actor_info        |  
| address            |  
| category          |  
| city               |  
| country            |  
| customer           |  
| customer_list     |  
| film               |  
| film_actor        |  
| ...                |
```

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

Quick I_S Example

```
mysql> select table_name, table_rows, avg_row_length
from tables where table_schema='sakila';
```

```
+-----+-----+-----+
| table_name          | table_rows | avg_row_length |
+-----+-----+-----+
| actor               |          200 |          25 |
| actor_info          |          NULL |          NULL |
| address             |          603 |          64 |
| category            |           16 |          20 |
| city                |          600 |          22 |
| country             |          109 |          22 |
| customer            |          599 |          70 |
| customer_list       |          NULL |          NULL |
| film                |         1000 |         135 |
| film_actor          |         5462 |           9 |
...

```

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

Index Selectivity Example

```
SELECT
  t.TABLE_SCHEMA
  , t.TABLE_NAME
  , s.INDEX_NAME
  , s.COLUMN_NAME
  , s.SEQ_IN_INDEX
  , (
    SELECT MAX(SEQ_IN_INDEX)
    FROM INFORMATION_SCHEMA.STATISTICS s2
    WHERE s.TABLE_SCHEMA = s2.TABLE_SCHEMA
    AND s.TABLE_NAME = s2.TABLE_NAME
    AND s.INDEX_NAME = s2.INDEX_NAME
  ) AS "COLS_IN_INDEX"
  , s.CARDINALITY AS "CARD"
  , t.TABLE_ROWS AS "ROWS"
  , ROUND((s.CARDINALITY / IFNULL(t.TABLE_ROWS, 0.01)) *
100, 2) AS "SEL %"
FROM INFORMATION_SCHEMA.STATISTICS s
  INNER JOIN INFORMATION_SCHEMA.TABLES t
    ON s.TABLE_SCHEMA = t.TABLE_SCHEMA
    AND s.TABLE_NAME = t.TABLE_NAME
WHERE t.TABLE_SCHEMA != 'mysql'
  AND t.TABLE_ROWS > 10
  AND s.CARDINALITY IS NOT NULL
  AND (s.CARDINALITY / IFNULL(t.TABLE_ROWS, 0.01)) < 1.00
ORDER BY t.TABLE_SCHEMA, t.TABLE_NAME, s.INDEX_NAME, "SEL %";
```

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

why add to I_S?

- need additional or different server or database meta-data
- want to make system information available for querying through the database
- `show` command doesn't cut it
- like to tinker (good introduction to MySQL source)

Our Need

- To see disk usage of the various MySQL databases

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

Setup

- MySQL source (currently 5.1.17)
- Ubuntu (running virtually on OS X)
- GCC and other compiler tools
- editor

5 steps to mod I_S

1. Add mnemonic
2. Add case for switch
3. Define column structure
4. Associate table with function
5. Create function to generate data

Step 1: add mnemonic

- `sql/table.h` (~line 515-550)
- In `enum_schema_tables` enumeration
- Add new mnemonic:

```
enum enum_schema_tables
{
    SCH_CHARSETS= 0,
    SCH_COLLATIONS,
    SCH_COLLATION_CHARACTER_SET_APPLICABILITY,
    SCH_COLUMNS,
    SCH_COLUMN_PRIVILEGES,
    SCH_DISKUSAGE,
    SCH_ENGINES,
    SCH_EVENTS,
    ...
}
```

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

Step 2: add case

- `sql/sql_parse.cc` (~line 1350-1450)
- In `prepare_schema_table` function
- Add case statement:

```
int prepare_schema_table(THD *thd, LEX *lex, Table_ident *table_ident,
                        enum enum_schema_tables schema_table_idx)
{
...
    case SCH_SCHEMATA:
        case SCH_DISKUSAGE:
    case SCH_ENGINES:
...
}
```

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

Step 3: create structure

- `sql/sql_show.cc` (~line 5400-5900)
- Create new `ST_FIELD_INFO` array:

```
ST_FIELD_INFO collation_fields_info[]=
{
...
};
```

```
ST_FIELD_INFO disk_usage_fields_info[]=
{
  {"DATABASE", 40, MYSQL_TYPE_STRING, 0, 0, "Database"},
  {"SIZE (Kb)", 21, MYSQL_TYPE_LONG, 0, 0, "Size (Kb)"},
  {0, 0, MYSQL_TYPE_STRING, 0, 0, 0}
};
```

```
ST_FIELD_INFO engines_fields_info[]=
{
...
};
```

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

Step 4: associate table

- `sql/sql_show.cc` (~line 5900-5975)
- Add table to `schema_tables` array:

```
...  
{ "COLUMN_PRIVILEGES", column_privileges_fields_info, create_schema_table,  
  fill_schema_column_privileges, 0, 0, -1, -1, 0 },  
  
{ "DISKUSAGE", disk_usage_fields_info, create_schema_table,  
  fill_disk_usage, make_old_format, 0, -1, -1, 0 },  
  
{ "ENGINES", engines_fields_info, create_schema_table,  
  fill_schema_engines, make_old_format, 0, -1, -1, 0 },  
...
```

- Must be ordered as `enum_schema_tables` (step 1)

Step 5: create function

- `sql/sql_show.cc` (~line 2800-5400)
- Create `fill_disk_usage` function

```
int fill_disk_usage(THD *thd, TABLE_LIST *tables, COND *cond)
{
    TABLE *table= tables->table;
    CHARSET_INFO *scs= system_charset_info;
    ...
    table->field[0]->store("InnoDB TableSpace", strlen("InnoDB TableSpace"), scs);
    table->field[1]->store((longlong)fsize, TRUE);
    if (schema_table_store_record(thd, table))
        DEBUG_RETURN(1);
    ...
}
```

Compile and Use

```
mysql> select table_name,table_schema from information_schema.tables;
+-----+-----+
| table_name                | table_schema |
+-----+-----+
| CHARACTER_SETS            | information_schema |
| COLLATIONS                | information_schema |
| COLLATION_CHARACTER_SET_APPLICABILITY | information_schema |
| COLUMNS                  | information_schema |
| COLUMN_PRIVILEGES        | information_schema |
| DISKUSAGE                 | information_schema |
| ENGINES                   | information_schema |
| EVENTS                    | information_schema |
| FILES                     | information_schema |
| ...                       |              |
| USER_PRIVILEGES          | information_schema |
| VIEWS                     | information_schema |
+-----+-----+
28 rows in set (0.00 sec)
```

Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

```

mysql> describe diskusage;
+-----+-----+-----+-----+-----+-----+
| Field      | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| DATABASE   | varchar(40)   | NO   |     |         |       |
| SIZE (Kb)  | bigint(21)    | NO   |     | 0       |       |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> select * from diskusage;
+-----+-----+
| DATABASE          | SIZE (Kb) |
+-----+-----+
| InnoDB TableSpace |          0 |
| InnoDB Logs       |          0 |
| BACKUP            |       16384 |
| mysql             |      690108 |
| sakila            |     3834192 |
| test              |         8192 |
+-----+-----+
6 rows in set (0.01 sec)

mysql> select * from diskusage where `database` rlike 'Inno';
+-----+-----+
| DATABASE          | SIZE (Kb) |
+-----+-----+
| InnoDB TableSpace |          0 |
| InnoDB Logs       |          0 |
+-----+-----+
2 rows in set (0.00 sec)

```

Creating INFORMATION_SCHEMA Tables
 Mike Kruckenberg
 MySQL Conference 2007

5 steps to mod I_S

1. Add mnemonic
2. Add case for switch
3. Define column structure
4. Associate table with function
5. Create function to generate data

Pluggable I_S

```
ST_FIELD_INFO vmstat_field_info[]=
{
...
};

static int fill_vmstat_schema(THD *thd, TABLE_LIST *tables, COND *cond)
{
...
}

static int vmstat_is_plugin_init(void *p)
{
...
}

static int vmstat_is_plugin_deinit(void *p)
{
...
}

mysql_declare_plugin(vmstat_is)
{
...
}
```



Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007

Other Thoughts

- I_S is a standard, might not be *right* to get too carried away with this
- self-compiled code add new issues
- migrate SHOW commands

More Info

- **Systems Table Article**
<http://www.dbazine.com/db2/db2-disarticles/pelzer4>
- **Pluggable I_S (Brian Aker)**
http://krow.net/index.pl?node_id=2994 http://krow.net/index.pl?node_id=3063
- **Two books:**



Creating INFORMATION_SCHEMA Tables
Mike Kruckenberg
MySQL Conference 2007