MySQL Workbench and InnoDB GIS

Yoshiaki Yamasaki / 山崎 由章

MySQL Senior Sales Consultant, Asia Pacific and Japan

ORACLE

MySQL

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

MySQL Workbench

MySQL Workbench

- MySQLの管理/開発支援ツール
- Windows 7 and newer (32 and 64bits)
- Mac (Snow Leopard and newer)
- Linux (Ubuntu, Fedora, Oracle Linux or sources)
- Latest version 6.2 **New**
- MySQL 5.1以降に対応

ORACLE

MySQL Workbench

- Free/Open Source Community Edition
- Commercial Edition

  GUI for Online Backup, Audit, DBDoc etc

  Included with commercial MySQL Editions

- http://dev.mysql.com/downloads/workbench/

MySQL Workbench

- 管理
- 開発支援
- データモデリング
- マイグレーション

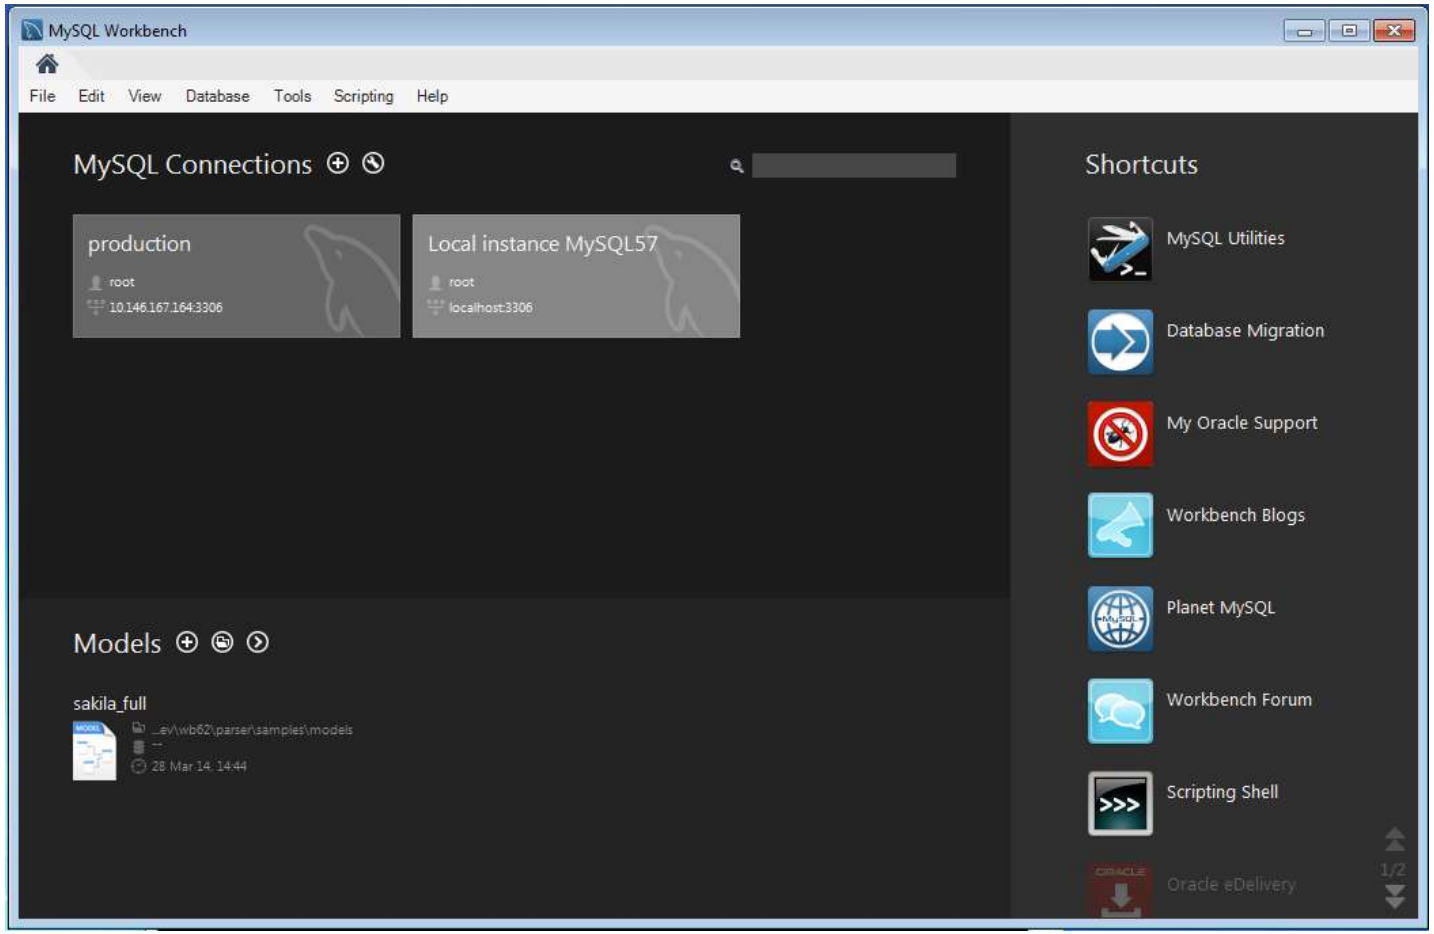

MySQL Workbench – 管理

- Server Status
- Start/Stop
- Connections
- Accounts
- Dump/Import
- Logs
- Configuration
- Performance Reports

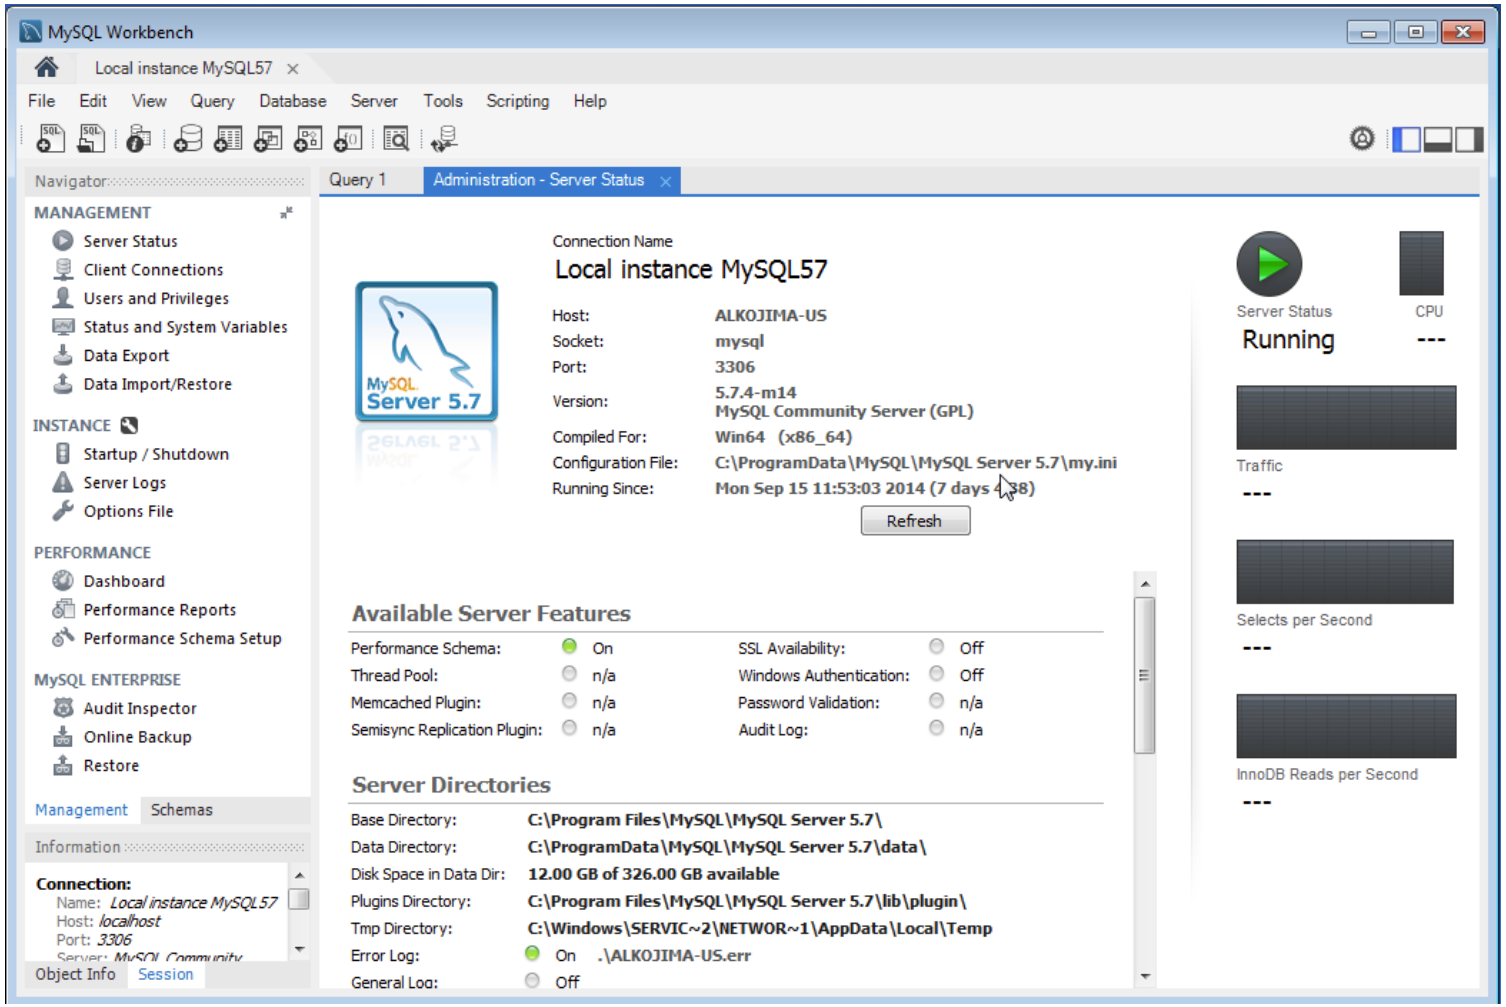

MySQL Workbench – 開発支援：SQL Editor

- Browse Databases
- View/Edit Data
- Run Queries and Scripts
- Create/Alter Objects (online)
- Query Optimization (Visual Explain etc)

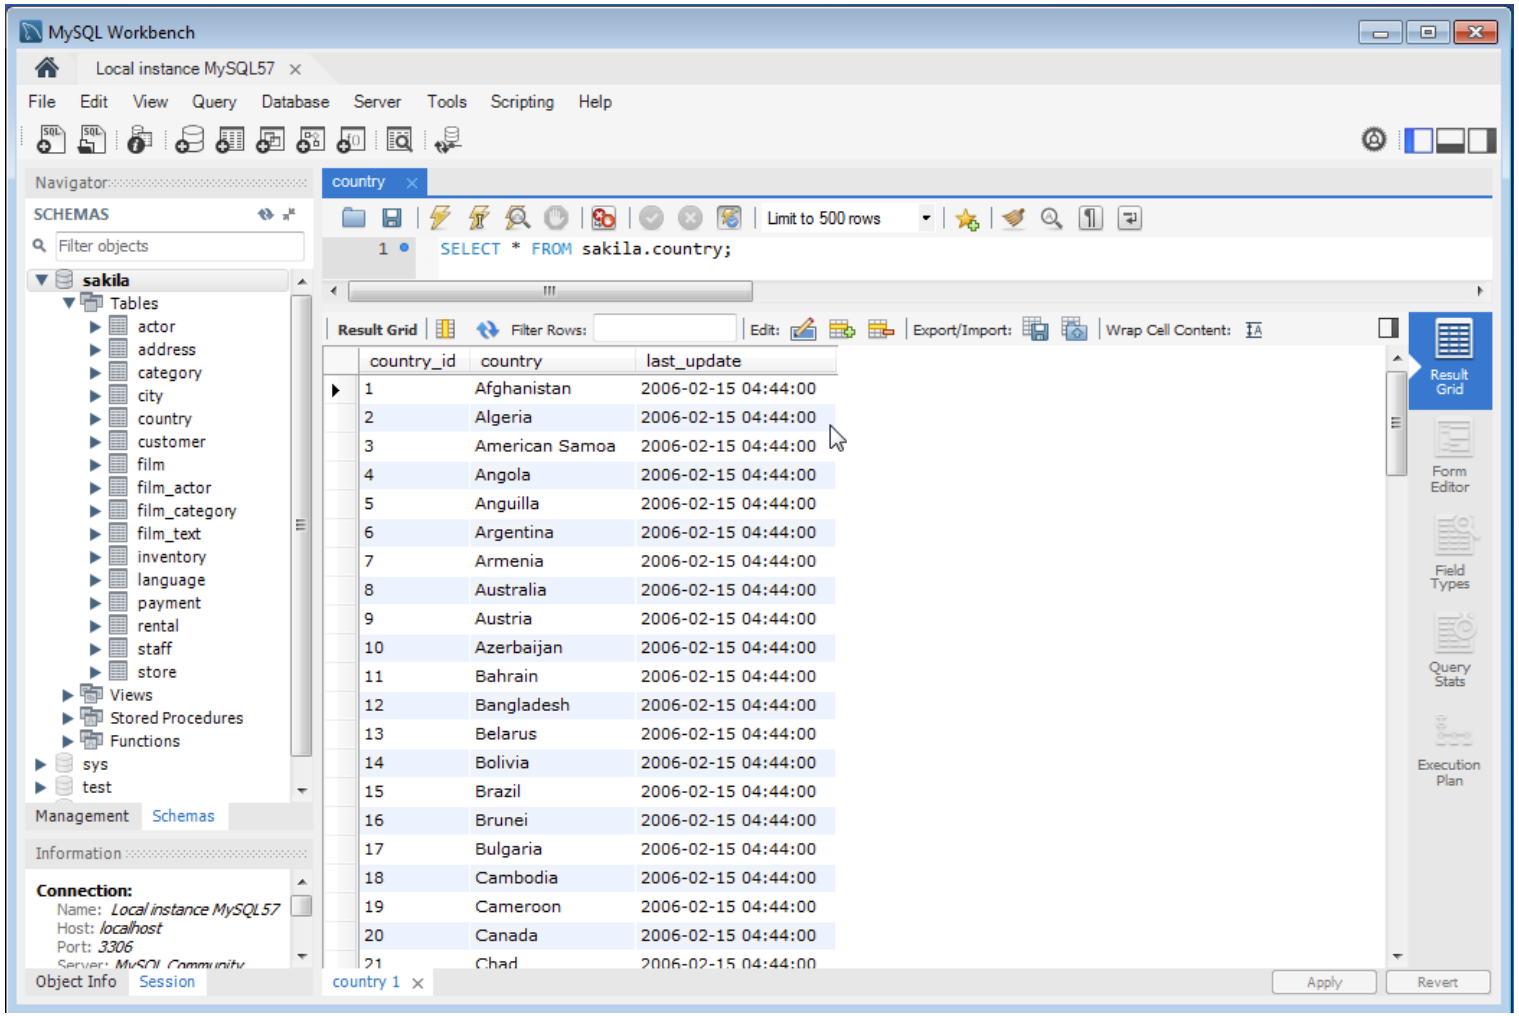

MySQL Workbench – モデリング

- Visually model database schemas
- Reverse engineer existing databases
- Generate SQL code from diagram
- Synchronize

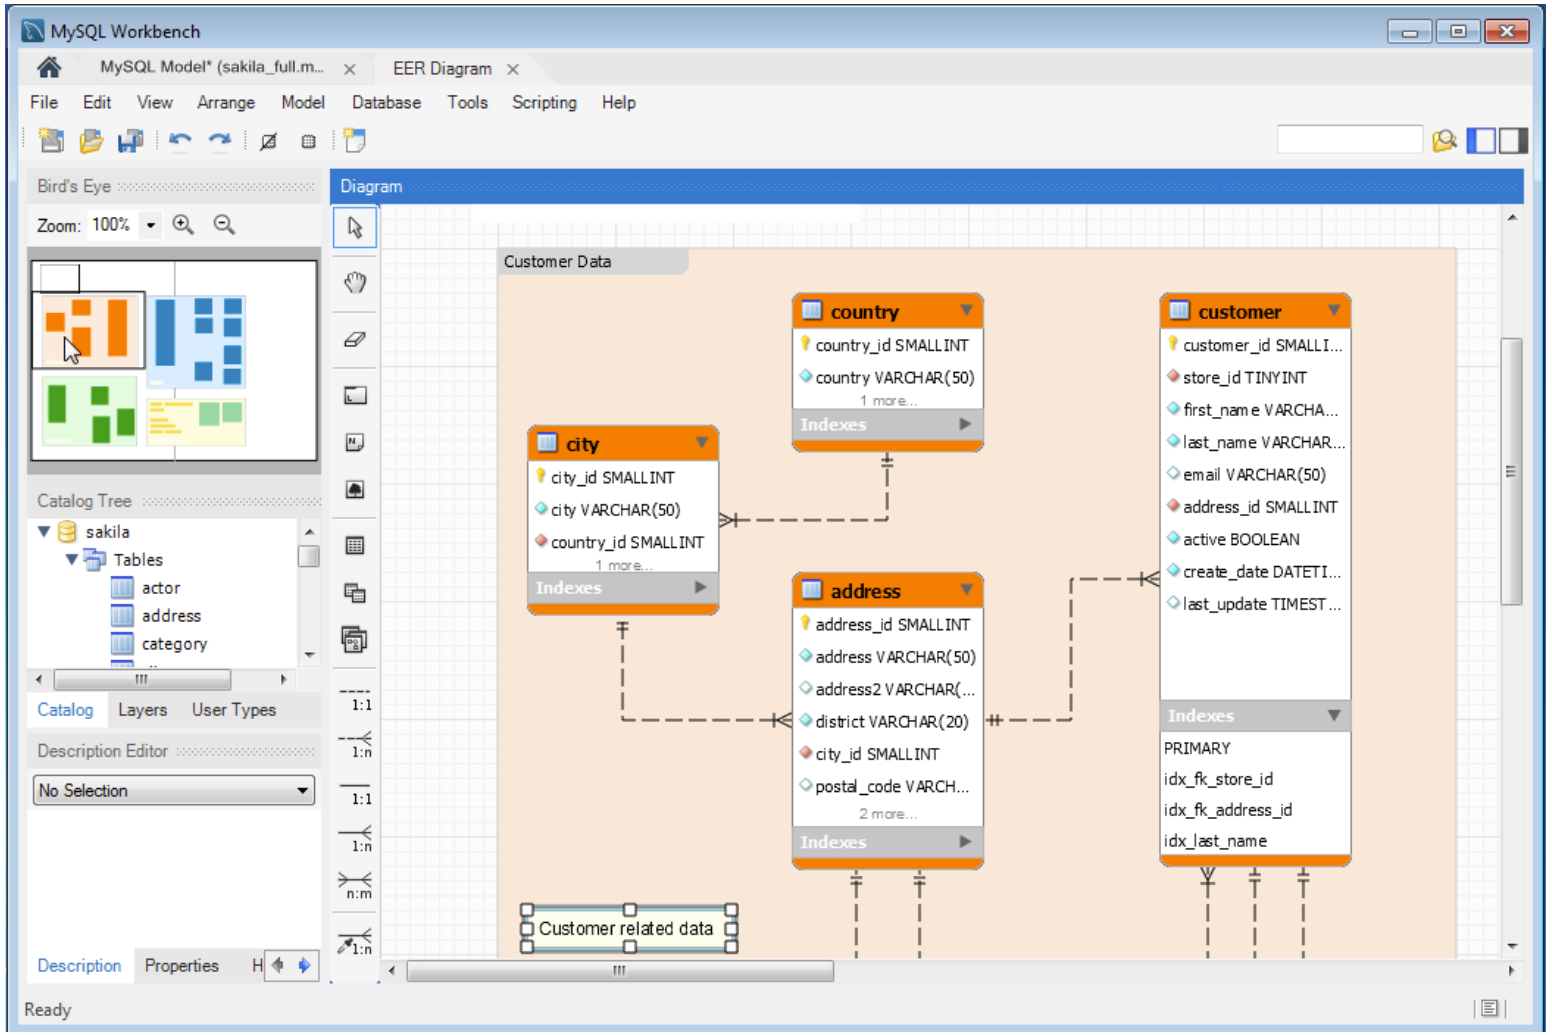

MySQL Workbench 6.2

GA

- Fabric対応
  - Fabricノードの追加、構成確認、接続
- Performance Dashboard
  - パフォーマンススキーマのレポートとグラフ
- Visual Explain
- GIS Viewer
- マイグレーション
  - New Microsoft Access
  - Microsoft SQL Server, Sybase, PostgreSQL

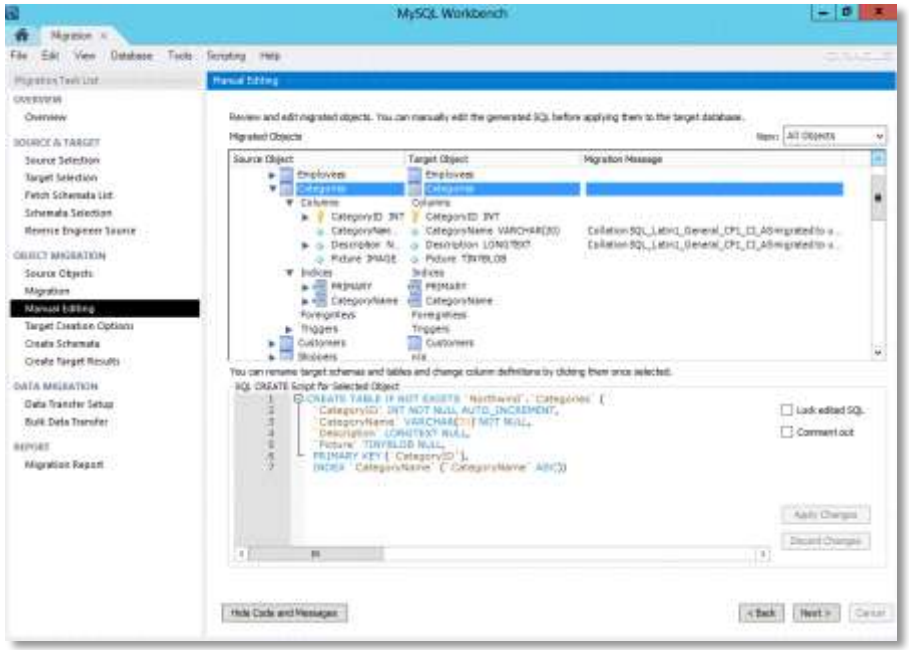

MySQL 5.7: Optimizer - JSON EXPLAINへのコスト情報追加

- JSON EXPLAINを拡張
  - 出力可能なコスト情報を全て表示
  - MySQL WorkbenchのVisual Explainにも表示

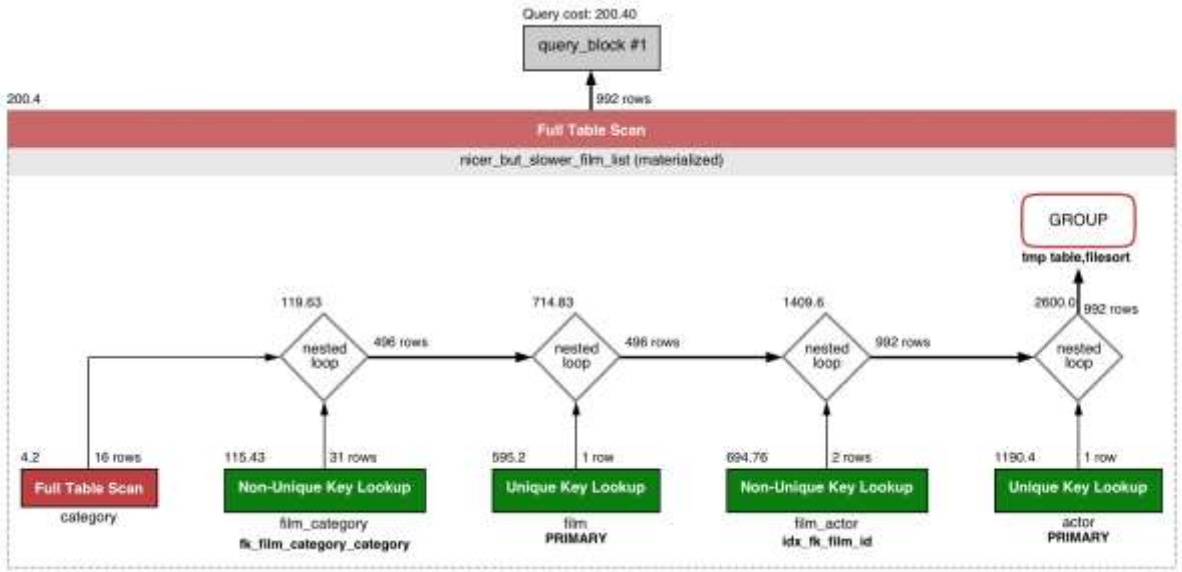


```
{
  "query_block": {
    "select_id": 1,
    "cost_info": {
      "query_cost": "200.40"
    },
    "table": {
      "table_name": "nicer_but_slower_film_list",
      "access_type": "ALL",
      "rows_examined_per_scan": 992,
      "rows_produced_per_join": 992,
      "filtered": 100,
      "cost_info": {
        "read_cost": "2.00",
        "eval_cost": "198.40",
        "prefix_cost": "200.40",
        "data_read_per_join": "852K"
      },
      "used_columns": [
        "FID",
        "title",
        "description",
        "category",
        "price",
        "length",
        "rating",
        "actors"
      ],
...
```

スキーマとデータをMySQLへマイグレーション

- MS SQL Server
  - 2000
  - 2005
  - 2008
  - 2012
- MS Access
- Sybase ASE
- Sybase SQL Anywhere

New in WB 6.2

- PostgreSQL
- SQLite
- Generic
- MySQL

MySQL Workbench 6.2 その他の改善点

- Query Results Panelの改善
- SQLスクリプトの実行
  - サイズの大きなSQLファイルを読み込まずに直接実行
- スニペットの共有
  - よく使うSQL文を共有して使用
- メタデータ・ロック・ブラウザー
- クライアント・コネクション・ブラウザーの改善
  - OS,プログラム,クライアントバージョン等の詳細情報を表示
- など

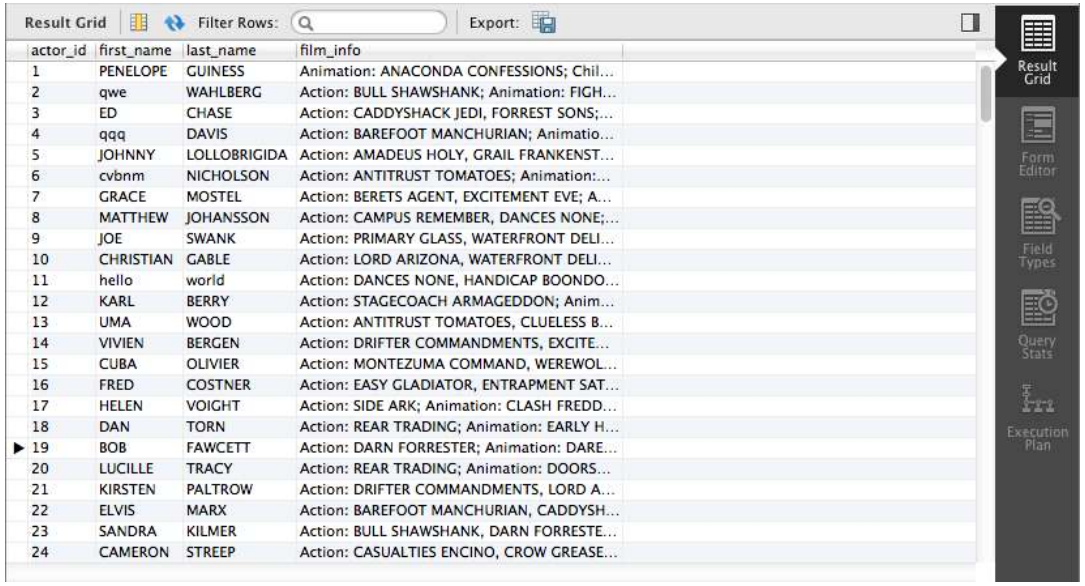

actor_id	first_name	last_name	film_info
1	PENELOPE	GUINESS	Animation: ANACONDA CONFESSIONS; Chil...
2	qwe	WAHLBERG	Action: BULL SHAWSHANK; Animation: FIGH...
3	ED	CHASE	Action: CADDYSHACK JEDI, FORREST SONS;...
4	qqq	DAVIS	Action: BAREFOOT MANCHURIAN; Animatio...
5	JOHNNY	LOLLOBRIGIDA	Action: AMADEUS HOLY, GRAIL FRANKENST...
6	cvbnm	NICHOLSON	Action: ANTITRUST TOMATOES; Animation:...
7	GRACE	MOSTEL	Action: BERETS AGENT, EXCITEMENT EVE; A...
8	MATTHEW	JOHANSSON	Action: CAMPUS REMEMBER, DANCES NONE;...
9	JOE	SWANK	Action: PRIMARY GLASS, WATERFRONT DELI...
10	CHRISTIAN	GABLE	Action: LORD ARIZONA, WATERFRONT DELI...
11	hello	world	Action: DANCES NONE, HANDICAP BOONDO...
12	KARL	BERRY	Action: STAGECOACH ARMAGEDDON; Anim...
13	UMA	WOOD	Action: ANTITRUST TOMATOES, CLUELESS B...
14	VIVIEN	BERGEN	Action: DRIFTER COMMANDMENTS, EXCITE...
15	CUBA	OLIVIER	Action: MONTEZUMA COMMAND, WEREWOL...
16	FRED	COSTNER	Action: EASY GLADIATOR, ENTRAPMENT SAT...
17	HELEN	VOIGHT	Action: SIDE ARK; Animation: CLASH FREDD...
18	DAN	TORN	Action: REAR TRADING; Animation: EARLY H...
▶ 19	BOB	FAWCETT	Action: DARN FORRESTER; Animation: DARE...
20	LUCILLE	TRACY	Action: REAR TRADING; Animation: DOORS...
21	KIRSTEN	PALTROW	Action: DRIFTER COMMANDMENTS, LORD A...
22	ELVIS	MARX	Action: BAREFOOT MANCHURIAN, CADDYSH...
23	SANDRA	KILMER	Action: BULL SHAWSHANK, DARN FORRESTE...
24	CAMERON	STREEP	Action: CASUALTIES ENCINO, CROW GREASE...

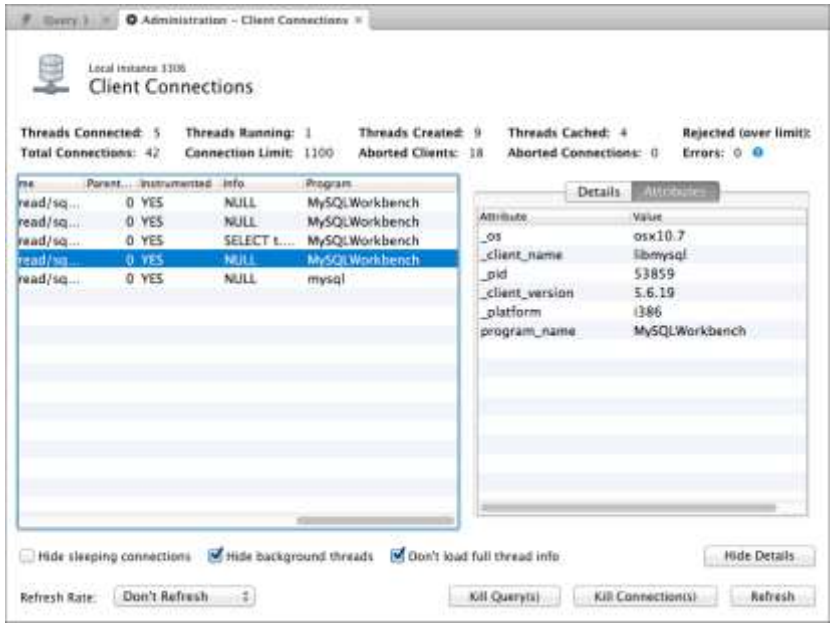

"Spatial Viewer" and "Geometry Viewer"

- New in Workbench 6.2

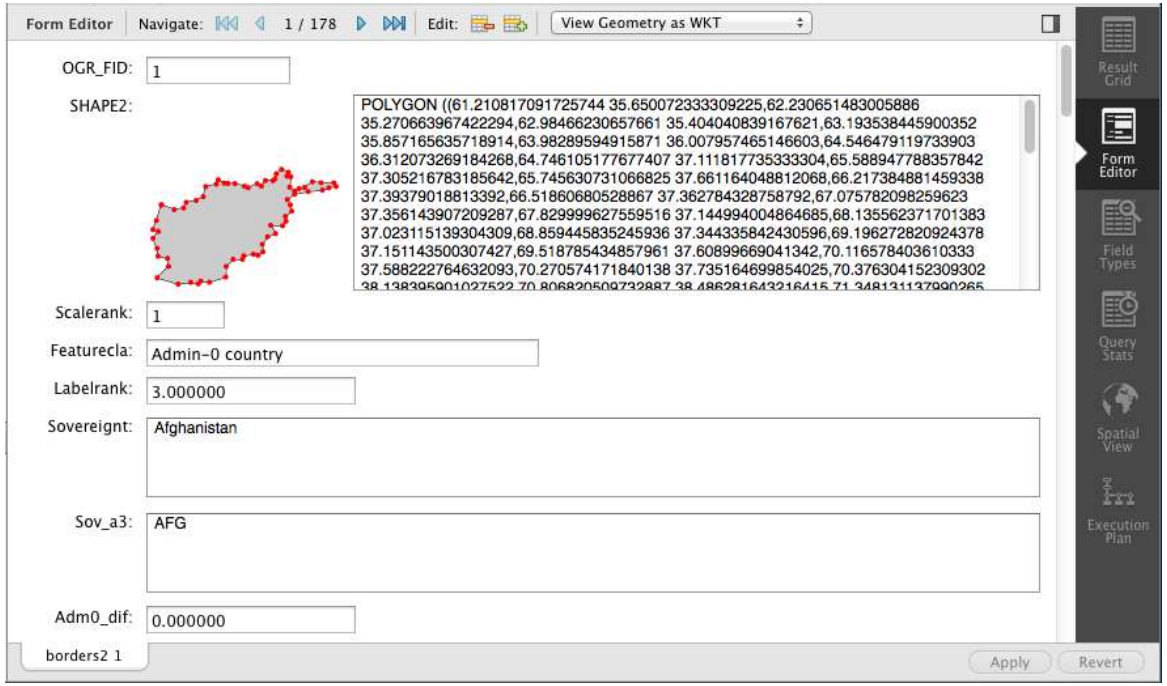

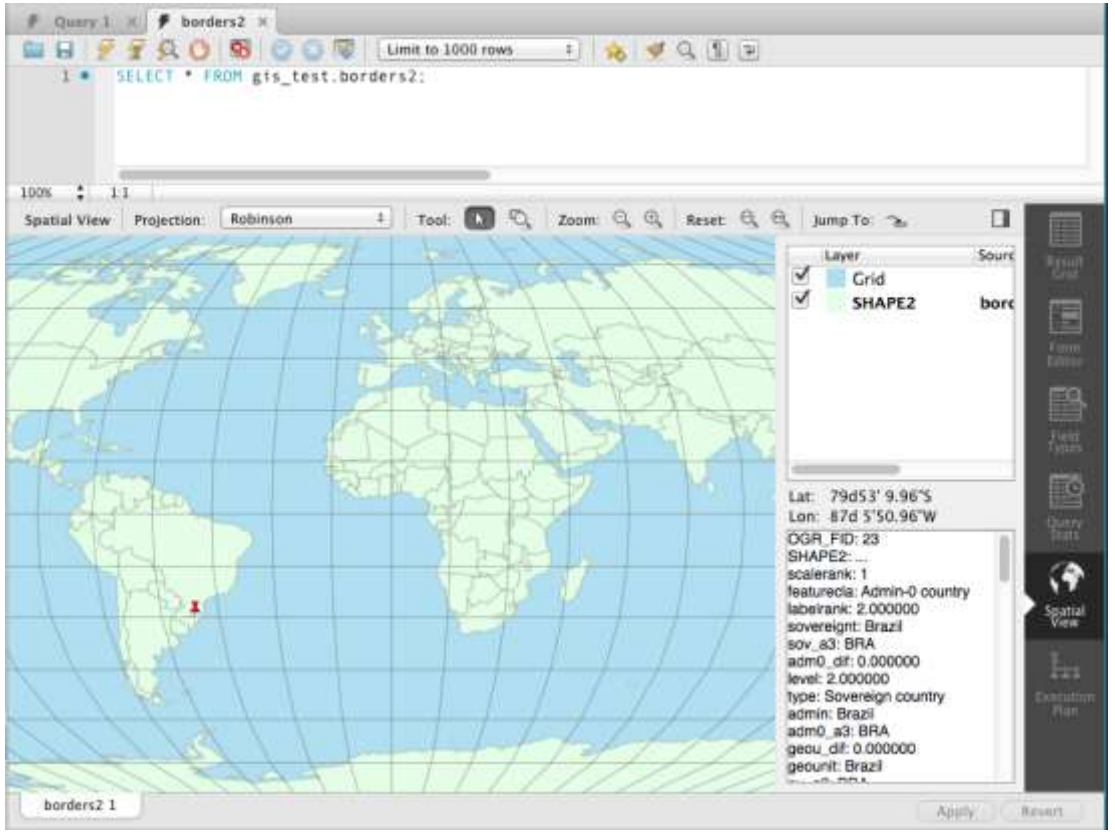

MySQL 5.7: GIS - Boost.Geometryとの統合

- 独自コードの置き換え
  - 空間図形情報の計算
  - 空間図形情報の分析
- OGC(Open Geospatial Consortium)準拠
  - パフォーマンスの向上
- Boost.Geometryによる効果
  - エキスパートとの交流
  - 非常に活発なコミュニティ
- Boost.Geometryへのコントリビュートも

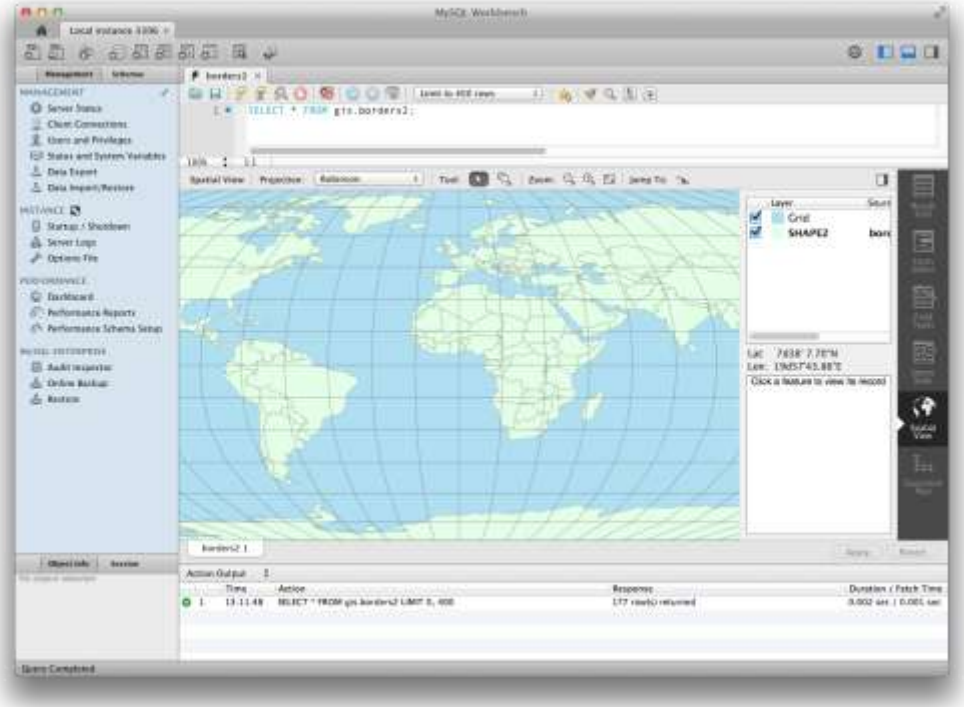

Common Terms and Concepts

Standards Organizations

- Open Geospatial Consortium (OGC)
  - Set and maintain the ISO SQL standards for GIS
    - Also many others: transformations, markup languages (KML, GML, etc.), presentation, …
- European Petroleum Survey Group (EPSG)
  - An authority for things such as coordinate reference systems
    - CRS/EPSGID/SRID
  - Now part of the OGP
- Environmental Systems Research Institute
  - A commercial company that is a de-facto standard
    - Creators of the very popular Shapefile (.shp) format
    - Creators of the very popular ArcGIS software

Common Terms

- Coordinates
  - x,y,z coordinates in planar space (4D is m or measure)
  - MySQL currently only supports 2D (x,y) coordinates
- Projection
  - Allows a spheroidal surface to be represented in planar format
  - Necessary for creating “flat” or 2D maps from locations on a spheroid
- Coordinate reference system (CRS/SRS/EPSGID/SRID)
  - Defines where a POINT—represented by a longitude and latitude coordinate pair—is located on the physical earth and defines its relationship to other POINTs
  - Also used for calculating distances

Data Formats

- Vector
  - Points, lines, and polygons
  - Position (x,y,z) is relative in a coordinate system
  - Generally used by database servers
  - Includes .Shp, .OSM, .KML, .GeoJSON, …
- Raster
  - Cells in a grid matrix, tied to an anchor (e.g. the {1,1} cell)
  - Generally used in aerial, satellite, and other imagery
  - Includes .tiff, .jpg, .gif, and other pixel based formats

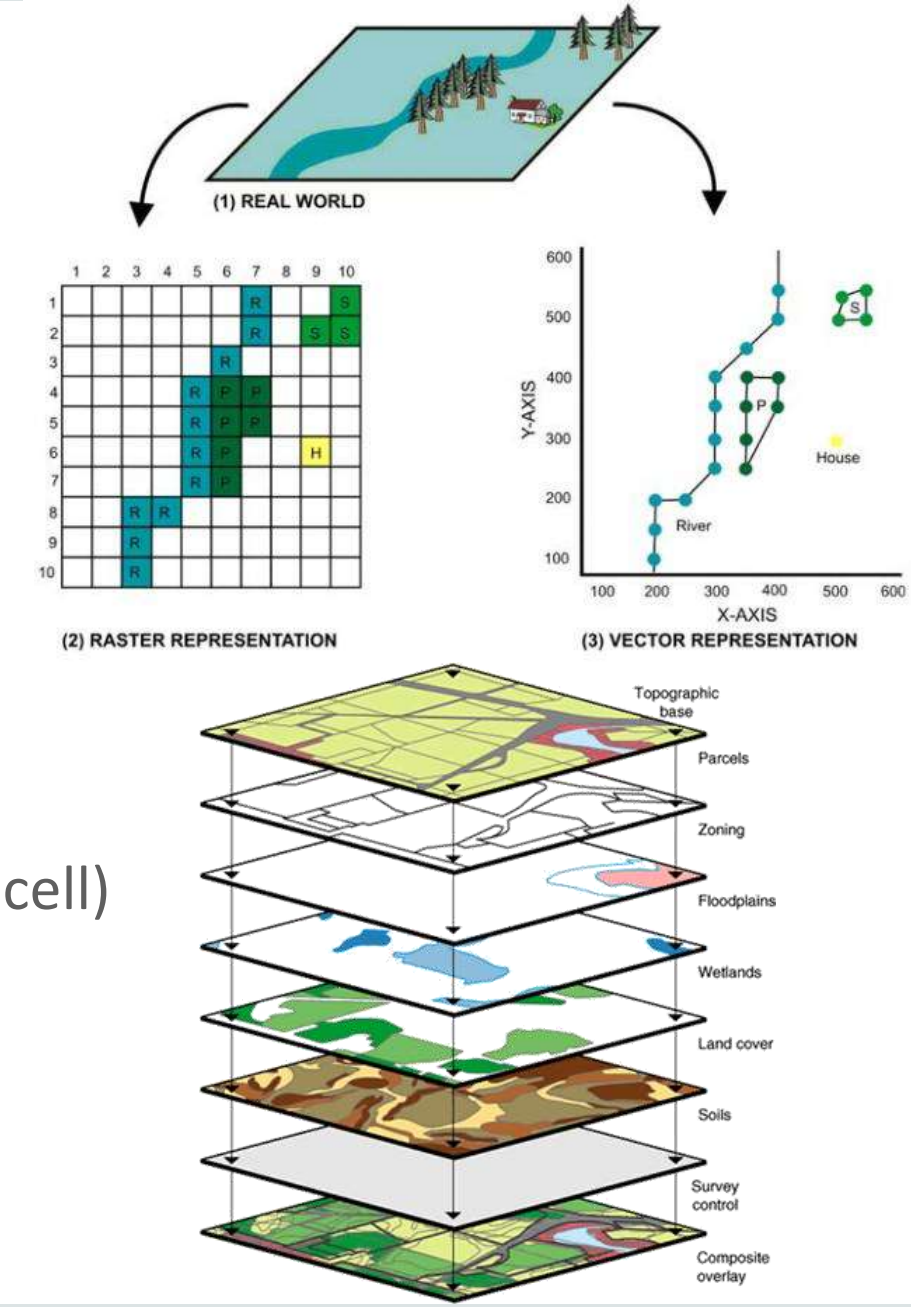

Data Sources

- Free
  - OpenStreetMap
  - Governments and NGOs
  - Universities (UCGIS) and other non-profits
- Commercial / Non-free
  - Data Depot, Geography Network, Land Info, Macon, NEXRAIN, SPOT image, …
- Custom
  - Geoencoding from various sources, such as user generated images and GPS data
    - Most media today is automatically geotagged: tweets, photos, Facebook posts, …
  - Create custom maps using ArcGIS, QGIS, GRASS, …

Migrating Data

- The OSGeo project
  - Geospatial Data Abstraction Library (GDAL/OGR)
    - Import data from various vector formats
    - Convert raster based data to vector format
- ESRI
  - ArcGIS
    - ArcSDE geodatabase abstraction layer for interfacing directly with database servers
    - Convert data between various file formats
- Open Street Map
  - Perl (OsmDB.pm) and Java (Osmosis) tools for importing OSM data

What's New in MySQL 5.7

ORACLE

Boost.Geometryとの統合

- 独自コードの置き換え
  - 空間図形情報の計算
  - 空間図形情報の分析
- OGC(Open Geospatial Consortium)準拠
  - パフォーマンスの向上
- Boost.Geometryによる効果
  - エキスパートとの交流
  - 非常に活発なコミュニティ
- Boost.Geometryへのコントリビュートも

Spatial Indexes for InnoDB

- R-tree based
  - トランザクションサポート
  - ファントムリードを防ぐロック
  - レコードは最小のbounding box (境界線)を含む
  - 現時点では2Dデータのみをサポート
    - 3Dデータのサポートも計画中
  - historical spatial index DDL をサポート

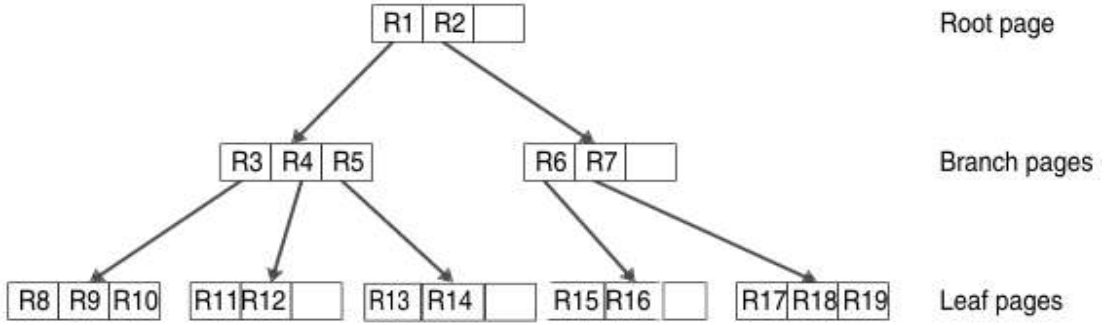

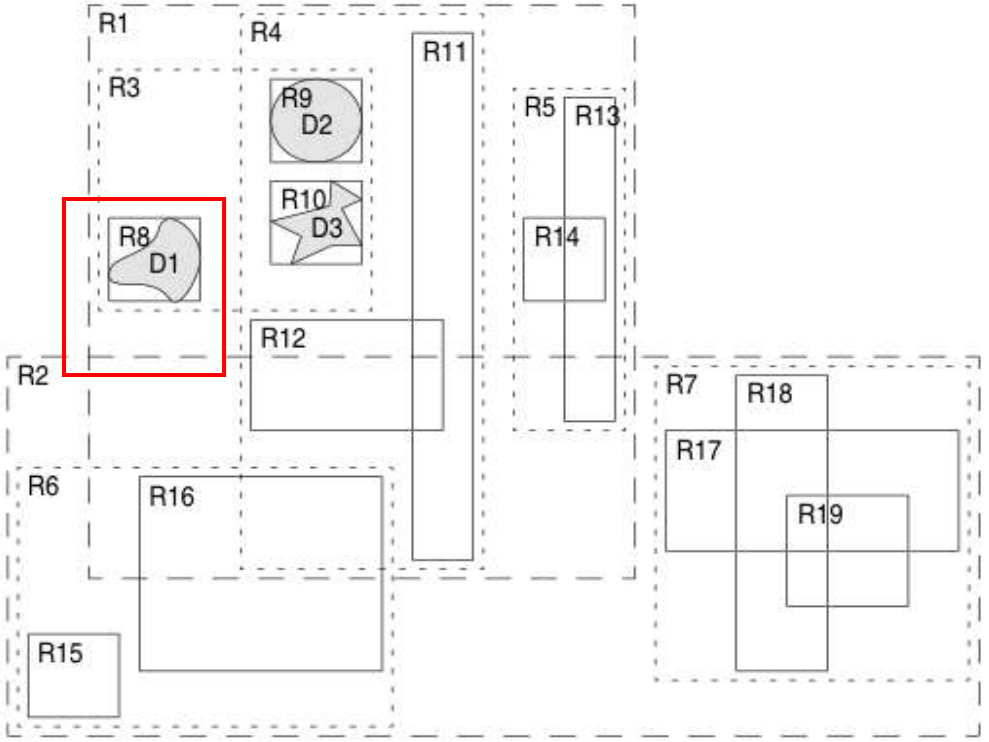

Additional Features

- GeoHash
  - B-tree indexes on the generated hash values
  - Quick lookups for exact matches
  - Not very accurate for proximity searches
- GeoJSON
- Additional functions
  - ST_IsValid(), ST_IsSimple(), …
- Limited SRID support

GeoJSON Example

```
{
  "type": "Feature",
  "geometry": {
    "type": "Point",
    "coordinates": [125.6, 10.1]
  },
  "properties": {
    "name": "Dinagat Islands"
  }
}
```

Some Real World Examples

A Starting Point

- My old apartment in Brooklyn, NY
  - 33 Withers Street Brooklyn, NY 11211
  - POINT(<LONG>,<LAT>)
    - -73.951353,40.716914

https://www.google.com/maps/place/33+Withers+St,+Brooklyn,+NY+11211/@40.7169144,-73.9513538

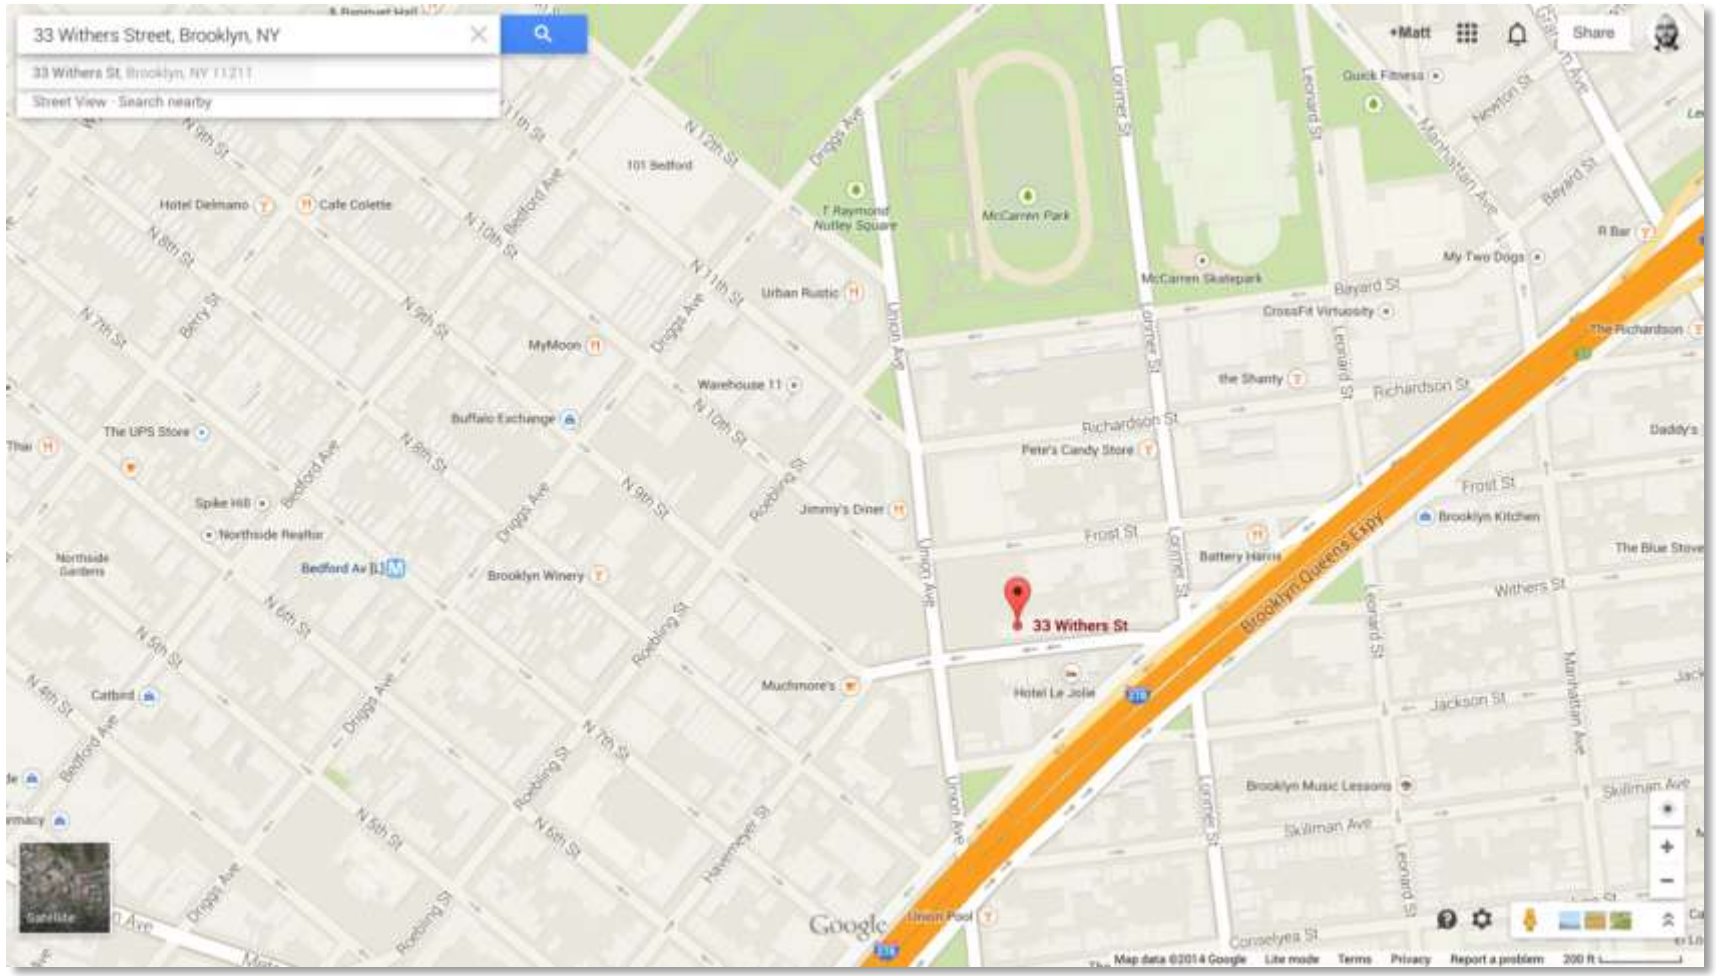

The Application Use Case

- I’m hungry and in the mood for Thai food
  - What Thai restaurants are around me?
  - What’s the closest one?
  - Can I see the menu, contact info, yelp ratings, etc.?
  - How would I get there?

Getting Some Data In

- Download a NYC OSM extract:
  - http://osm-extracted-metros.s3.amazonaws.com/new-york.osm.bz2
- Import the data using a customized OsmDB.pm Perl module
  - http://wiki.openstreetmap.org/wiki/OsmDB.pm (original)
  - https://www.dropbox.com/s/l17vj3wf9y13tee/osmdb-scripts.tar.gz (customized)
    - Creates a Geometry column named 'geom'
    - Creates a spatial index on the 'geom' column

```
mysql -e "create database nyosm"
bunzip2 new-york.osm.bz2
./bulkDB.pl new-york.osm nyosm
```

The Generated Schema

- http://wiki.openstreetmap.org/wiki/Elements

```
mysql> show tables;
+-----------------+
| Tables_in_nyosm |
+-----------------+
| nodes           |
| nodetags        |
| relationmembers |
| relations       |
| relationtags    |
| waynodes        |
| ways            |
| waytags         |
+-----------------+
```

- We'll focus on nodes and nodetags for our examples
- Nodes
  - A point or location
- Nodetags
  - Metadata about each location
  - X name/value pairs

De-normalizing the Tag Data

- Greatly simplify our query
- Allow for the use of a full-text index
  - Also improves performance
- Mimic better schema created by osm2pgsql
  - http://wiki.openstreetmap.org/wiki/Osm2pgsql/schema#planet_osm_nodes

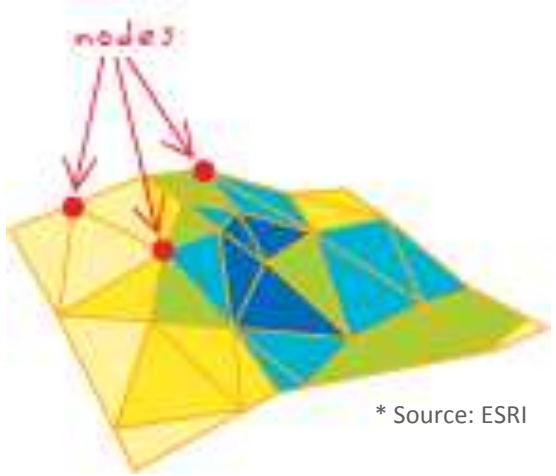

* Source: ESRI

```
mysql> alter table nodes add column tags text, add fulltext index(tags);
mysql> update nodes set tags=(SELECT group_concat(concat(k, "=", v) SEPARATOR
';') from nodetags where nodetags.id=nodes.id group by nodes.id);
```

Final Nodes Table

```
mysql> show create table nodes¥G
*************************** 1. row ***************************
       Table: nodes
Create Table: CREATE TABLE `nodes` (
  `id` bigint(20) DEFAULT NULL,
  `geom` geometry NOT NULL,
  `user` varchar(50) DEFAULT NULL,
  `version` int(11) DEFAULT NULL,
  `timestamp` varchar(20) DEFAULT NULL,
  `uid` int(11) DEFAULT NULL,
  `changeset` int(11) DEFAULT NULL,
  `tags` text,
  UNIQUE KEY `i_nodeids` (`id`),
  SPATIAL KEY `i_geomidx` (`geom`),
  FULLTEXT KEY `tags` (`tags`)
) ENGINE=InnoDB DEFAULT CHARSET=latin1
```

Creating a Distance Calculation Function

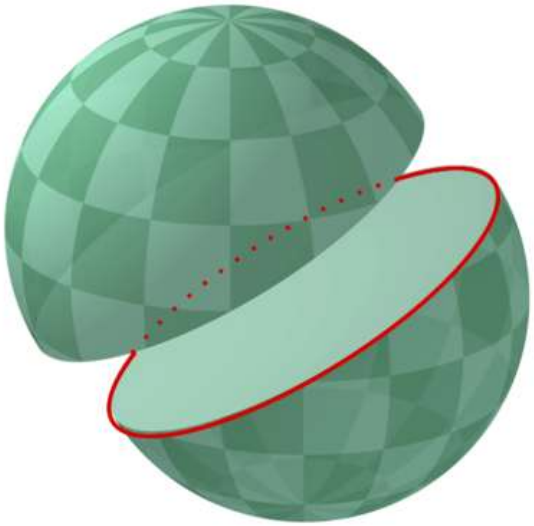

- Various great circle (orthodrome) distance formulas
  - Haversine, Spherical Law of Cosines (my choice), …
  - http://en.wikipedia.org/wiki/Great-circle_distance
  - Necessary for calculating distances between two Geometries
    - Need goes away when we support Geography and/or Projections (ST_Transform)

```
mysql> CREATE FUNCTION slc (lat1 double, lon1 double, lat2 double, lon2 double)
RETURNS double
RETURN 6371 * acos(cos(radians(lat1)) * cos(radians(lat2)) * cos(radians(lon2)
- radians(lon1)) + sin(radians(lat1)) * sin(radians(lat2)));
```

Creating a Bounding Box For Our Search

- Utilize the r-tree index by limiting area
  - Easy with future spatial reference systems support
    - WGS84 or SRID 4326 being the most common
  - Need to use some additional geographic formulas
    - http://www.movable-type.co.uk/scripts/latlong.html
    - Need should go away with full SRID support

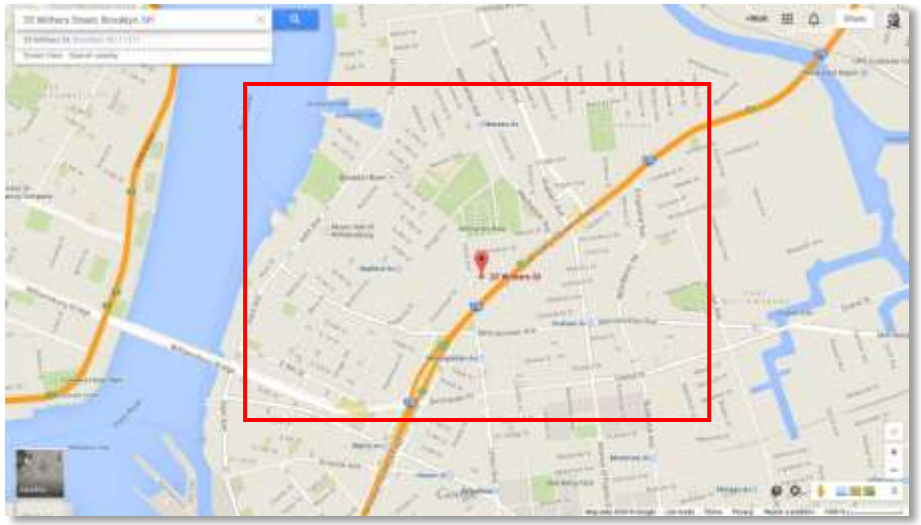

```
${origlon} = -73.951368
${origlat} = 40.716743
${lon1} = ${origlon} + (${distance_in_km}/abs(cos(radians({$origlat}))*111))
${lat1} = ${origlat} + (${distance_in_km}/111)
${lon2} = ${origlon} - (${distance_in_km}/abs(cos(radians({$origlat}))*111))
${lat2} = ${origlat} - (${distance_in_km}/111)
```

Calculating the Results

- Our final query, searching within ~ 10km radius

```
mysql> SELECT id,
slc(40.716743, -73.951368, y(geom), x(geom))*1000 as distance_in_meters,
tags, ST_AsText(geom)
FROM nodes
WHERE ST_Contains(ST_Envelope(linestring(point((-73.951368+(10/111)),
(40.716743+(10/111))), point((-73.951368-(10/111)), (40.716743-(10/111))))),
geom)
AND match(tags) against ("+thai +restaurant" IN BOOLEAN MODE)
ORDER BY distance_in_meters¥G
```

Examining the Results

```
*************************** 1. row ***************************
                id: 888976948
distance_in_meters: 614.4973960877276
              tags: addr:street=Bedford Avenue;amenity=restaurant;name=Tai
Thai;addr:housenumber=206;phone=7185995556
      astext(geom): POINT(-73.958637 40.717174)
*************************** 2. row ***************************
                id: 2178443635
distance_in_meters: 2780.87697408101
              tags: microbrewery=no;website=http://www.onemorethai.net/;name=One
More Thai;amenity=restaurant;opening_hours=12:00-22:30;cuisine=thai;phone=(212)
228-8858
      astext(geom): POINT(-73.983871 40.7210541)
*************************** 3. row ***************************

…
```

Mapping the Results

- From my old place
  - -73.951353,40.716914
- To Tai Thai
  - -73.958637,40.717174
- Maps APIs
  - Google, Bing, Apple, …

https://www.google.com/maps/dir/40.716914,+-73.951353/40.717174,+-73.958637

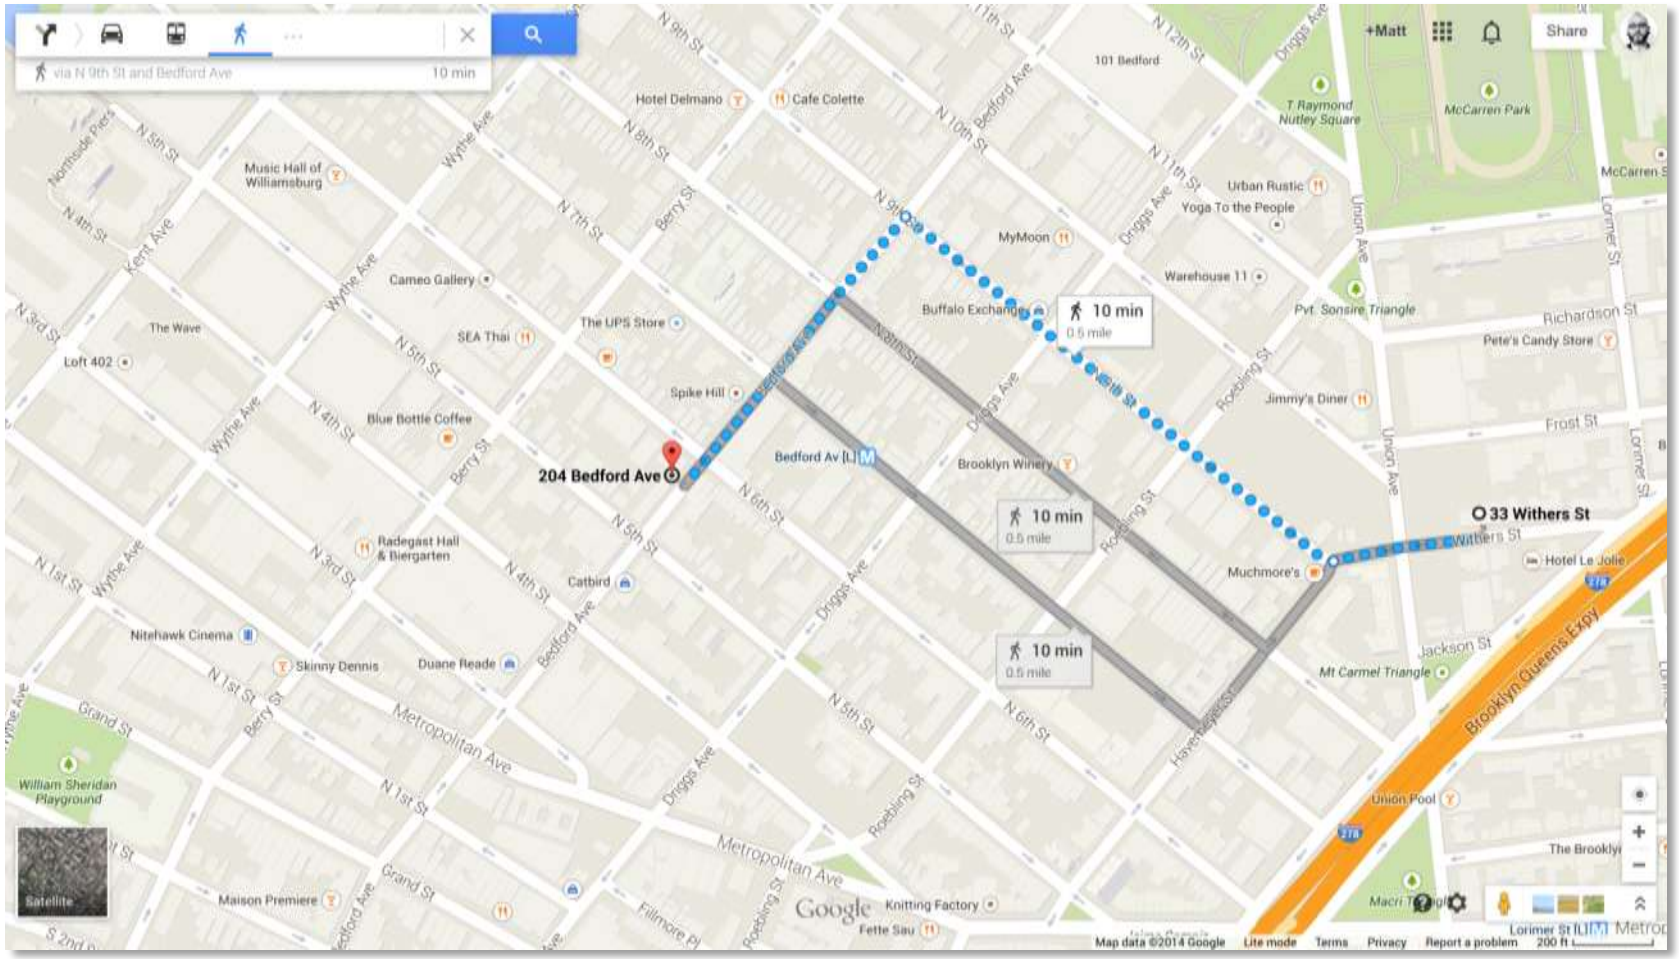

参考情報

- MySQL 5.7 and GIS, an Example
  http://mysqlserverteam.com/mysql-5-7-and-gis-an-example/
- Importing Raster Based Spatial Data into MySQL 5.7
  http://mysqlserverteam.com/importing-raster-based-spatial-data-into-mysql-5-7/
- MySQL GIS: Boost.Geometry
  - http://mysqlserverteam.com/making-use-of-boost-geometry-in-mysql-gis/
  - http://mysqlserverteam.com/why-boost-geometry-in-mysql/
  - http://mysqlserverteam.com/building-mysql-with-boost/
- MySQL GIS: InnoDB R-Tree
  - http://mysqlserverteam.com/innodb-spatial-indexes-in-5-7-4-lab-release/

ORACLE

What's Next for MySQL GIS

Storage Enhancements

- R-tree enhancements
  - 3D support
- Improved storage
  - Fixed length storage when possible
  - Transparent compression
  - Improved BLOB handling
- Concurrency improvements

Geography

- Geography types
- Geography functions
- Makes distance calculations very accurate
  - Simple ST_Distance() call for value in meters

Additional Features

- Projections
  - ST_Transform()
- 3D and Geodetic support
- OGC standard Information_Schema metadata
- Spatial reference system support
  - Starting with WGS84 (SRID 4326)
- Additional performance optimizations
- What else would you like to see?
  - Let us know!

Appendix : Additional Resources

- Manual
  - http://dev.mysql.com/doc/refman/5.7/en/spatial-extensions.html
- Community forum
  - http://forums.mysql.com/list.php?23
- Boost.Geometry
  - http://www.boost.org/libs/geometry
- Report GIS bugs and submit feature requests
  - http://bugs.mysql.com/