ORACLE

May 21, 2014
Walter E. Washington Convention Center
Washington, DC USA

Effectively Utilize LIDAR Data In Your Business Process

Daniel Geringer
Senior Software Development Manager

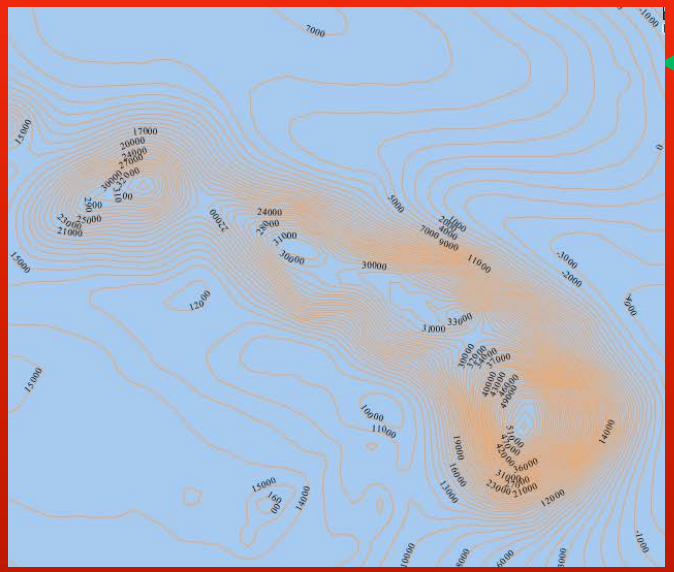

Program Agenda

- What is LIDAR Data
- Major Challenges
- LIDAR Data Model Optimized For Engineered Systems
- Loading / Compressing LIDAR Data
- Querying LIDAR DATA
- Derivative Products

What is LIDAR Data?

What is LIDAR Data

Laser Imaging Detection and Ranging (LIDAR) Data

- Optical remote sensors that collect millions of 3D points per second, along with numeric attributes associated with each point.
- Sensors targeted at the ground to generate precise elevation data models, sometimes called point clouds.
- Provides accurate representation of:
  - Railway infrastructure
  - Highways, roads, buildings, bridges
  - Forestry terrain
  - Bathymetry (sea floor elevations) and more…

LIDAR – Maximize Your Investment

- Costly to capture, but worth it.
- Accuracy (centimeter precision)
- Collection performed by:
  - Airplanes
  - Helicopters
  - Ground based mobile units
- Don't let your costly investment sit on a shelf

Perform Valuable Cost Saving Analysis

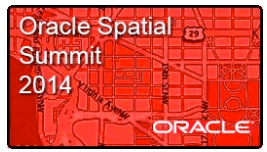

- Generate more accurate flood plain boundaries
- Forest / Coastal change detection for conservation and development purposes
  - High cost to survey remote areas for construction analysis
  - LIDAR data used to dynamically generate contours and DEM rasters
  - Reduce costs by not sending surveyors to remote areas for analysis
  - Help identify dunes, wetlands, valleys, protected flood plains

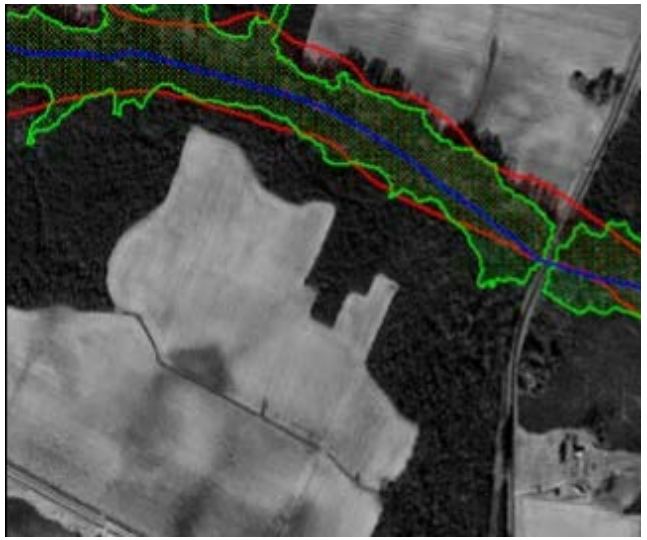

Major Challenges

LIDAR Data – Major Challenges

- Storage
  - Where do I put all this data? Archive?
  - Compress?
  - Optimal format for analysis? LAS, CSV, Proprietary
- Analysis
  - Derivative product generation (TINs, Contours, DEMs)
  - Spatial queries

Oracle 12c Spatial and Graph LIDAR Data Model

Optimized For Oracle Engineered Systems

What Is the Oracle Exadata Database Machine?

- Oracle SUN hardware uniquely engineered to work together with Oracle database software
- X4-2 Key features:
  - Database Grid – Up to 192 Intel cores connected by 40 Gb/second InfiniBand fabric, for massive parallel query processing.
  - Raw Disk – Up to 300 TB of uncompressed storage (high performance or high capacity)
  - Memory – Up to 4 TB
  - Hybrid Columnar Compression (HCC) – Query and archive modes available. 3x to 30x compression.
  - Storage Servers – Up to 14 storage servers (168 Intel cores) that can perform massive parallel smart scans. Smart scans offloads SQL predicate filtering to the raw data blocks. Results in much less data transferred, and dramatically improved performance.
  - Flash memory – Up to 44 TB

Exadata Configurations

- X4-2
  - Eight Rack – 24 cores
  - Quarter Rack – 48 cores
  - Half Rack – 96 cores
  - Full Rack -192 cores

Exadata X4-2 Quarter Rack Diagram

Exadata DB Machine X4-2 Quarter Rack

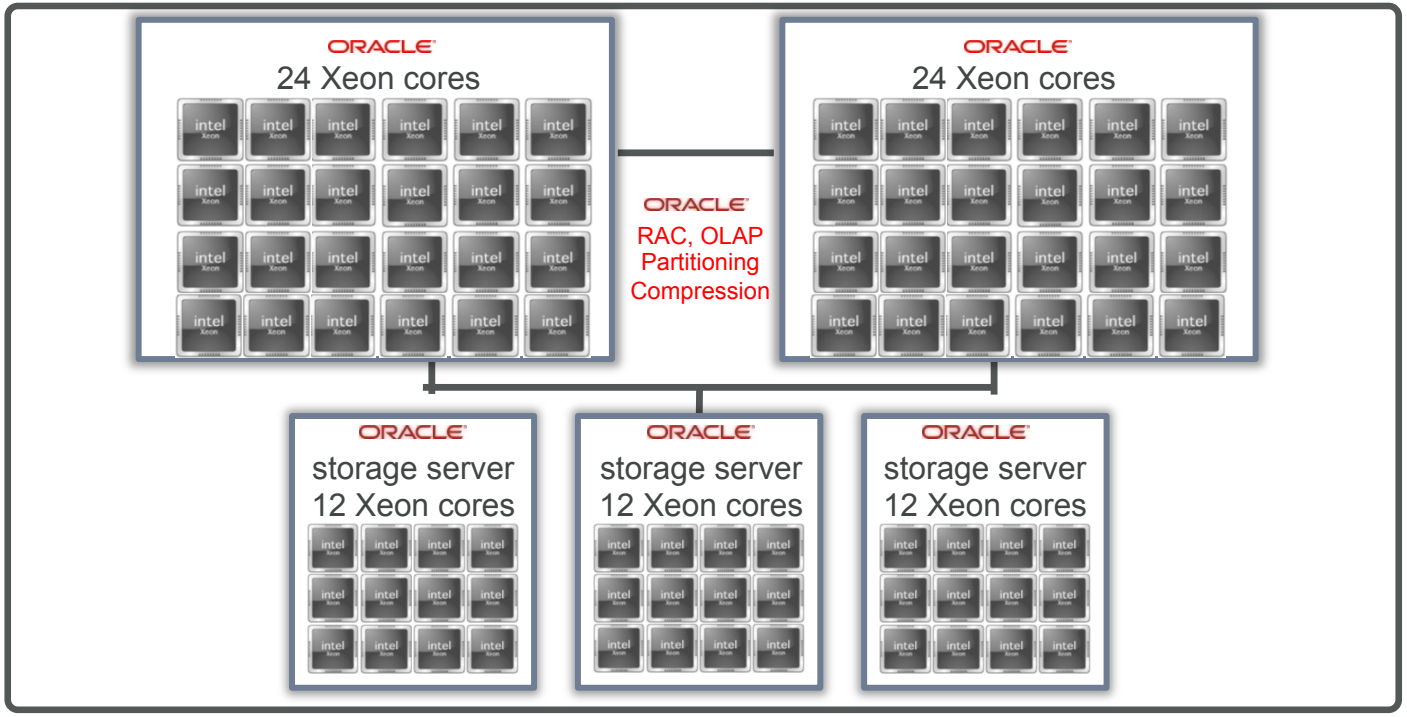

Oracle's Newest LIDAR Data Model

Production in December 2013

- "New paradigm" for LIDAR data, optimized for Oracle engineered systems.
- On Exadata and Sparc Supercluster (SSC), leverages Oracle's unique:
  - HCC compression technique, for extremely high compression rates
  - Parallel Enabled Smart Scan for extreme performance, including spatial queries.
  - **Spatial queries with no spatial/non-spatial indexes,** against LIDAR data in compressed form

LIDAR – Ingest and Compression Rates

Compression Used	Rows	Size GB	Compression Ratio
No Compression	**2,853,027,995**	**285.80**	--
Query Low	2,853,027,995	36.9	7.74 X
Query High	2,853,027,995	12.67	22.55 X
Archive Low	2,853,027,995	12.65	22.59 X
Archive High	2,853,027,995	9.28	30.79 X

Oracle's Newest LIDAR Data Model

(continued…)

- "Back to basics… a simple flat relational model
- X,Y,Z and attributes stored as ordinary Oracle NUMBER columns
- Table can contain other columns data types too.
- Works with any point table/view, not just LIDAR data

Loading LIDAR Data

Loading LIDAR – Two Common Formats

- CSV File
  - Create external table (points to CSV file)
  - Parallel CTAS (Create Table As Select)
- LAS File
  - New utility to create external table script (discovers LAS contents for you)
  - External table definition includes a “preprocessor” to stream LAS file
  - Create external table
  - Parallel CTAS

CSV File Load – External Table Example


```
CREATE TABLE  lidar_csv_ext (
  x           NUMBER,  y     NUMBER,  z       NUMBER,  intensity  NUMBER,
  returnval   NUMBER,  red   NUMBER,  green   NUMBER ,  blue       NUMBER)
organization external  (
  type oracle_loader
  default directory las_data_dir
  access parameters (
    records delimited by newline
    fields terminated by ",")
  location (
    -- More than one dat can be listed here. LAS filenames are comma delimited.
    'lidar_data.dat'  ))
reject limit unlimited;
```

LAS File Load – External Table Example

New Utility – Discovers LAS Columns and Builds Script


```
CREATE TABLE  lidar_las_ext (
  x          NUMBER,  y     NUMBER,  z       NUMBER,  intensity  NUMBER,
  returnval  NUMBER,  red   NUMBER,  green   NUMBER ,  blue       NUMBER,  pyramid_level  NUMBER )
organization external   (
  type oracle_loader
  default directory las_data_dir
  access parameters (
    records delimited by newline
    preprocessor sdo_pc_exe_dir:'preprocessor.sh'
    fields terminated by ",")
  location (
    -- More than one dat can be listed here. LAS filenames are comma delimited.
    'lidar_data_1.las',  'lidar_data_2.las' )  reject limit unlimited;
```

- Generates Pyramids
- Streams LAS File

Create Physical Table From External Table

- CTAS - Same strategy for CSV and LAS external tables
- Specify compression
- Select only columns you are interested in

```
CREATE TABLE las_data
  NOLOGGING
  COMPRESS FOR QUERY HIGH
  PARALLEL 96
  AS SELECT x,y,z,intensity,returnval,red,green,blue,pyramid_level
       FROM lidar_las_ext;
```

Pyramid – 3,265,110,000 Rows


```
SELECT pyramid_level, count(*)  FROM las_data  GROUP BY pyramid_level
ORDER BY pyramid_level;

PYRAMID_LEVEL            COUNT(*)
-----------------------  ---------------
              1          2938599000
              2          293860000
              3          29386000
              4          2939000
              5          294000
              6          29000
              7          3000
7 rows selected.
Elapsed: 00:00:03.55
```

LIDAR – Ingest Times – X4-2 Half RAC

(Query High – 22.5x Compression) (Query Low – 7.74x Compression)

Compression	CSV File 2,835,027,995 Rows	1000 LAS Files - Parallel 96 3,265,110,000 Rows	CTAS 2,835,027,995 Rows
Query Low	1 min 20 sec 35,662,849 rows/sec	5 min 27 sec 9,985,045 rows/sec	1 min 3 sec 44,641,339 rows/sec
Query High	2 min 2 sec 23,385,475 rows/sec	6 min 41 sec 9,547,105 rows/sec	1 min 46 sec 26,839,397 rows/sec
Archive Low	2 min 3 sec 23,008,290 rows/sec	7 min 49 sec 6,961,855 rows/sec	1 min 49 sec 26,107,503 rows/sec
Archive High	2 min 26 sec 19,408,353 rows/sec	14 min 36 sec 3,723,044 rows/sec	4 min 8 sec 11,467,154 rows/sec

Spatial Queries

Lidar Data For Haiti

Data Set Tested

- Each cell represents a LAS file
- Each cell contains approximately 2.9 million rows
- 2,853,027,995 total rows

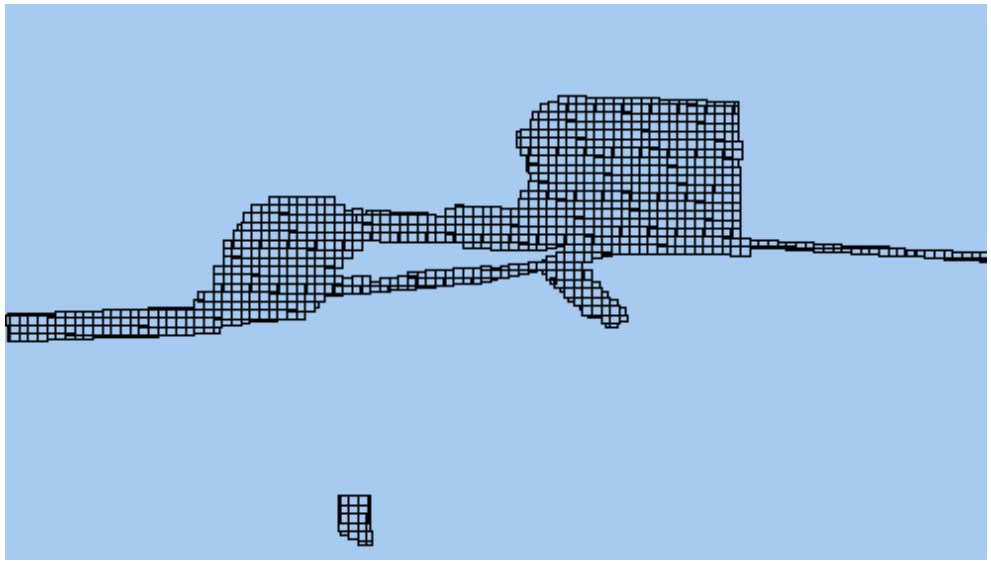

Box Queries

- Queries returned between 100 thousand and 100 million rows

Query ID	Rows Returned	Query Low Rows/Sec	Query High Rows/Sec	Achieve Low Rows/Sec	Archive High Rows/Sec
1	100,234	167,056	200,468	167,056	69,606
2	101,914	175,713	207,987	181,989	50,452
3	1,080,314	448,263	620,870	529,565	145,398
4	107,318	77,207	111,789	88,692	26,111
5	1,013,301	773,512	993,432	888,860	226,689
6	1,044,341	470,423	678,143	593,375	150,048
7	10,053,844	2,010,768	2,047,626	1,918,672	1,404,168
8	10,085,246	1,985,284	2,049,846	1,600,832	1,199,196
9	101,757,599	2,217,424	2,233,485	1,989,395	2,139,112

Any Polygon Shape

- Spatial queries are not limited to box.
- Any polygon shape. Perform analysis along:
  - a 300 meter buffer along a windy portion of a river
  - a 500 meter buffer along the location of a proposed road
  - on the fly aggregation of LIDAR data in environmentally protected polygonal areas

Polygon Query Example

- 300 meter buffer around a 9.44 km diagonal line
- Each cell contains approximately 3 million points
- Query returned 20,122,526 points
- 1,166,523 points/second
- Compressed For Query High

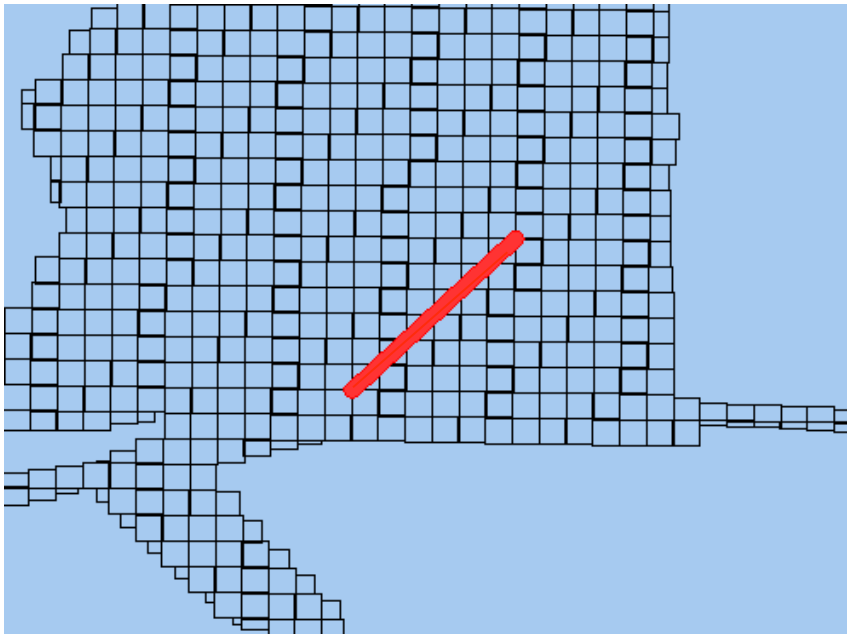

SPATIAL SEARCH IN COMPRESSED FORM

OVER 2,000,000 POINTS/SECOND

NO INDEX

Derivative Products

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.

Contours And DEM Rasters

Directly From LIDAR Model For Engineered Systems

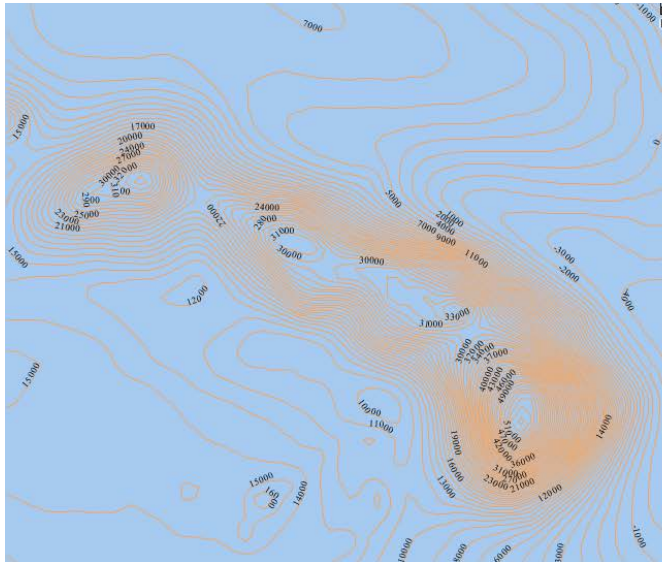

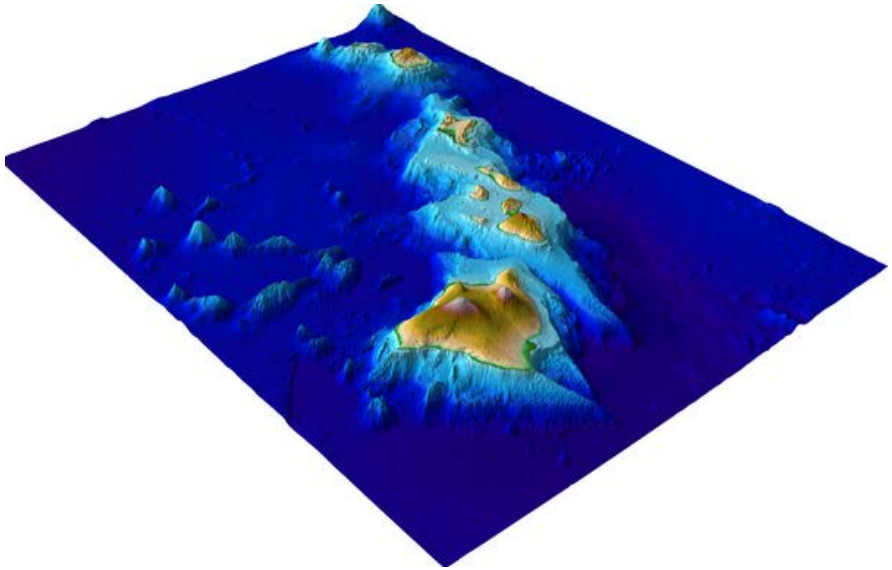

Summary - Key Differentiators

- High compression rates with HCC – Saves Storage!
- Searchable in compressed form – Smart Scan
- SQL – NO INDEX
  - Spatial Searches, no spatial or non-spatial indexes – Saves Storage!
  - Any polygon shape (not just box)
- Simple data model
- Simple to administer
  - Automatic Storage Management (ASM)
  - Don’t focus on distributing the data
  - Focus on your real requirement instead

ORACLE

May 21, 2014
Walter E. Washington Convention Center
Washington, DC USA

- **Next:**
  Lunch – Awards & SIG Presentation in Room 150AB
  (11:45am-1:15pm)

- **4:30-5:00pm: Meet the Experts – roundtable Q&A on topics in Room 150AB**
  - **Spatial Performance**
  - **Upgrading/Testing Apps for Spatial 12c**
  - **Raster & 3D**
  - **MapViewer/BI**
  - **Certification**
  - **SIG User Group**
- **Closing Reception (5:00) – Exhibit Hall**

Resources

Oracle Technology Network

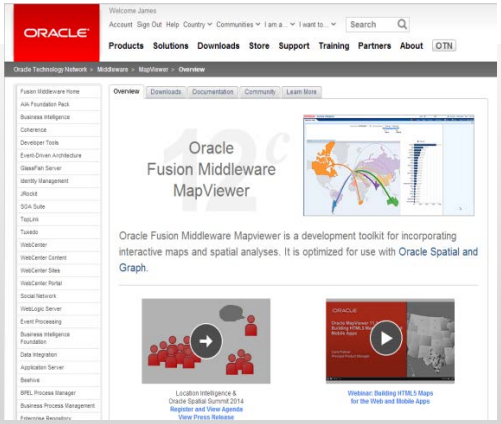

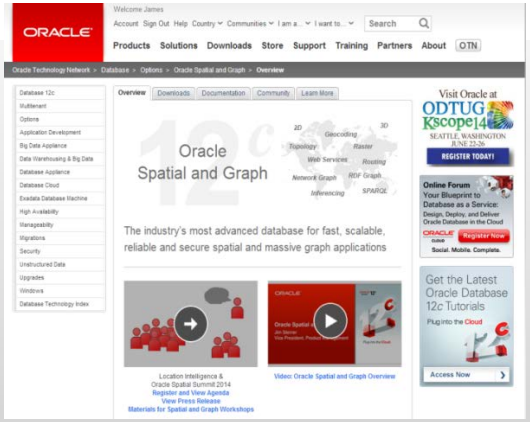

- www.oracle.com/technetwork/database/options/spatialand graph
- www.oracle.com/technetwork/middleware/mapviewer
- https://blogs.oracle.com ➔ oraclespatial
  ➔ oracle_maps_blog

Oracle Spatial & Graph Special Interest Group

Connect and exchange knowledge with the community of Spatial & Graph users

- **Talk with the Board this week**
  - Wednesday lunch – SIG Board presentation (150AB)
  - Stop by the SIG User Group roundtable at Meet the Experts, 4:30pm Wednesday in150AB
  - Visit Oracle's exhibitor table at breaks & sign up

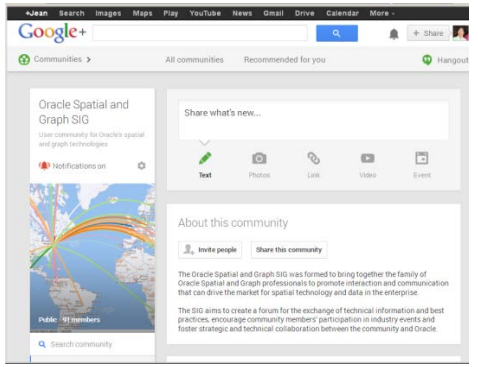

- **Join us**
  - Online communities: LinkedIn , Google+ , IOUG SIG (free membership)
  - Visit OTN Spatial Community page www.oracle.com/technetwork/database/options/spatialandgraph/community (or search online for **"Oracle Spatial and Graph Community"**)
  - Email oraclespatialsig@gmail.com

Spatial Certification & Partner Specialization

Get valuable credentials – differentiate your skills

- **Learn more at the Summit**
  - Wed, Track C 3:30 – Exam preparation session
  - Talk to us at Oracle's exhibitor table & "Meet the Experts" Certification table (Wed 4:30-5:00)
- **Take the next steps**
  - Schedule an exam, access topic lists / online training, learn about Partner Specialization requirements
    www.oracle.com/technetwork/database/options/spatialandgraph/learnmore/spatial-specialization-1835642.html
  - Online training materials for Certified Implementation Specialist exam
    https://competencycenter.oracle.com/opncc/full_glp.cc?group_id=22003