WMS Benchmarking 2011

Cadcorp GeognoSIS, Constellation-SDI,
GeoServer, Mapnik, MapServer,
QGIS Server

Executive summary

- Compare the performance of WMS servers
  - 6 teams
- In a number of different workloads:
  - Vector: projected (Google Mercator – EPSG:3857) street level
  - Raster: EPSG:4326 DEMs projected (Google Mercator – EPSG:3857)
- Data backends:
  - Vector: PostGIS
  - Raster: BIL

Benchmarking History

- [5th] FOSS4G benchmarking exercise. Past exercises included:
  - FOSS4G 2007: Refractions Research run - published the first comparison with the help of GeoServer and MapServer developers. Focus on big shapefiles, postgis, minimal styling (Brock Anderson & Justin Deoliveira)
  - FOSS4G 2008: OpenGeo run - published the second comparison with some review from the MapServer developers. Focus on simple thematic mapping, raster data access, WFS and tile caching (Justin Deoliveira & Andrea Aime)
  - FOSS4G 2009: MapServer and GeoServer teams in a cooperative benchmarking exercise (Andrea Aime & Jeff McKenna)
- Friendly competition: goal is to improve all software

Benchmarking 2010
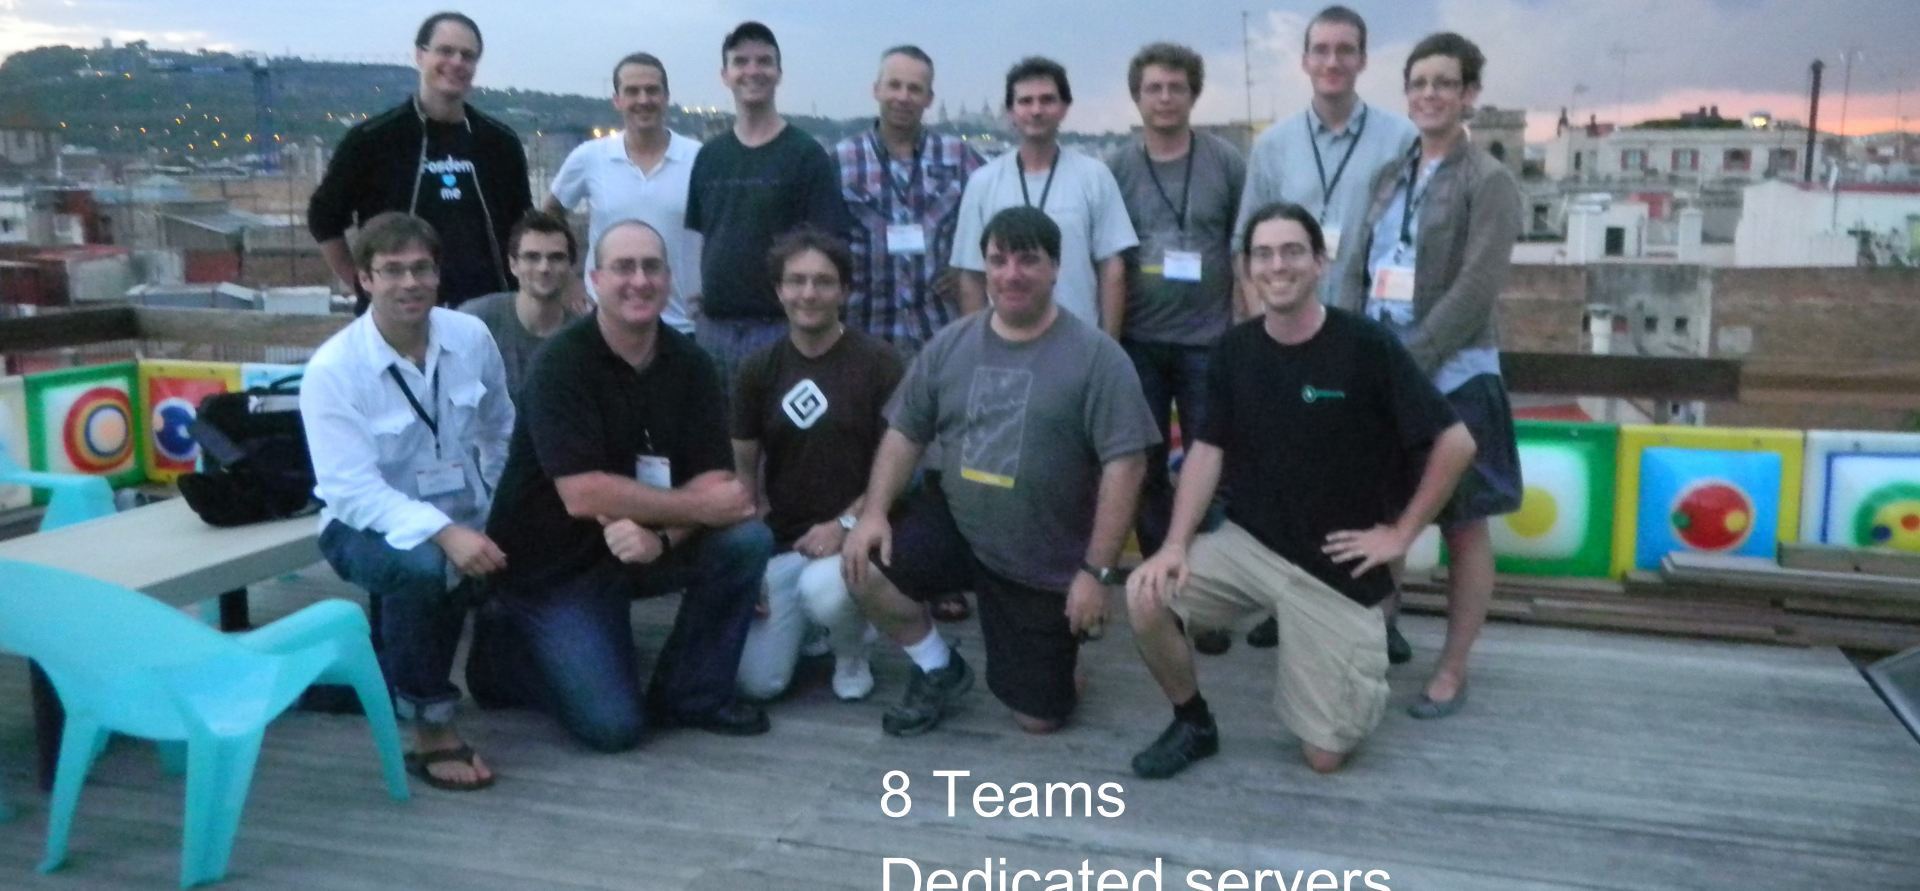
8 Teams
Dedicated servers
Area specific data set (Spain)

Benchmarking 2011

6 Teams
Dedicated hardware
Area specific dataset – Colorado

Rules of engagement

- Each server is tested in its latest version
- Each server performs exactly the same workload
  - Same set of WMS requests
  - Same data backends
  - Same image output format
- All modifications made to improve performance are to be included in a future release of the software
- Data used cannot be modified for display, other than indexing
- All testing to be done on the same benchmarking machines
  - Windows and Linux servers, 2 separate identical servers
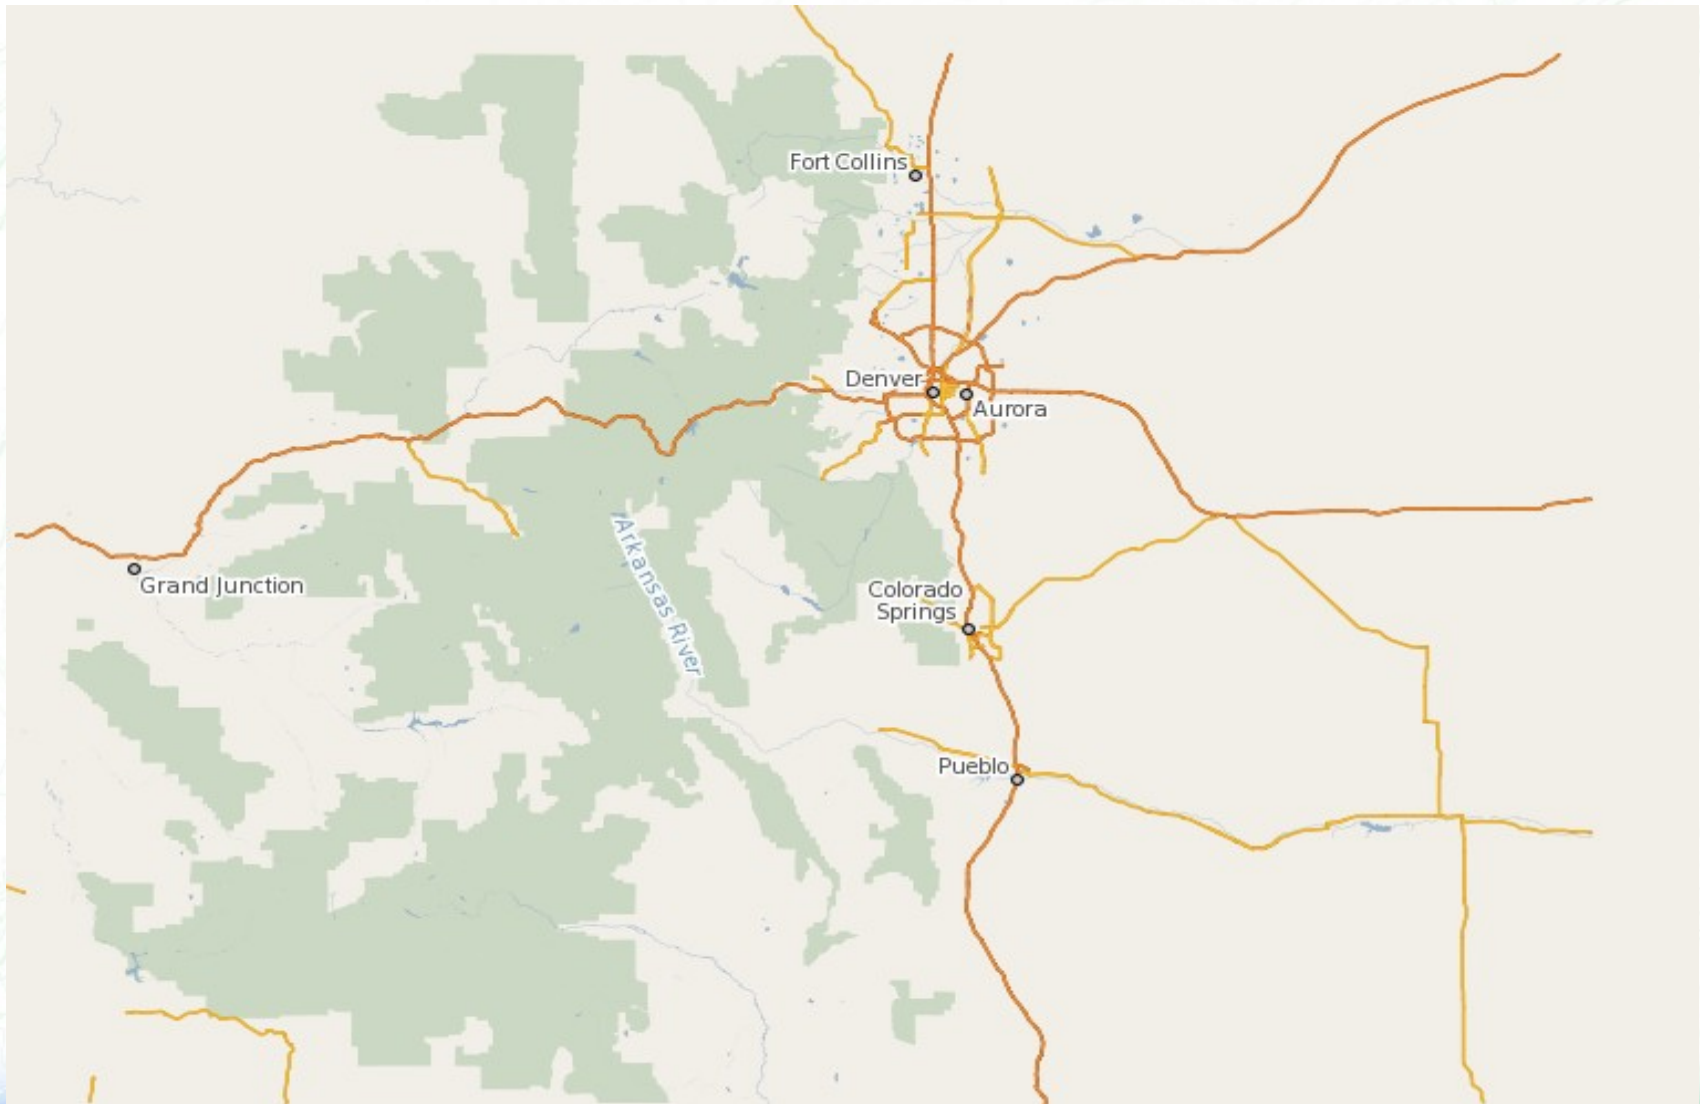
Datasets Used: Vector

Open Street Map for Colorado
- IMPOSM to import data to PostGIS
  - Optimized for rendering
- Styling from MapServer Utils Imposm branch (Googly style)
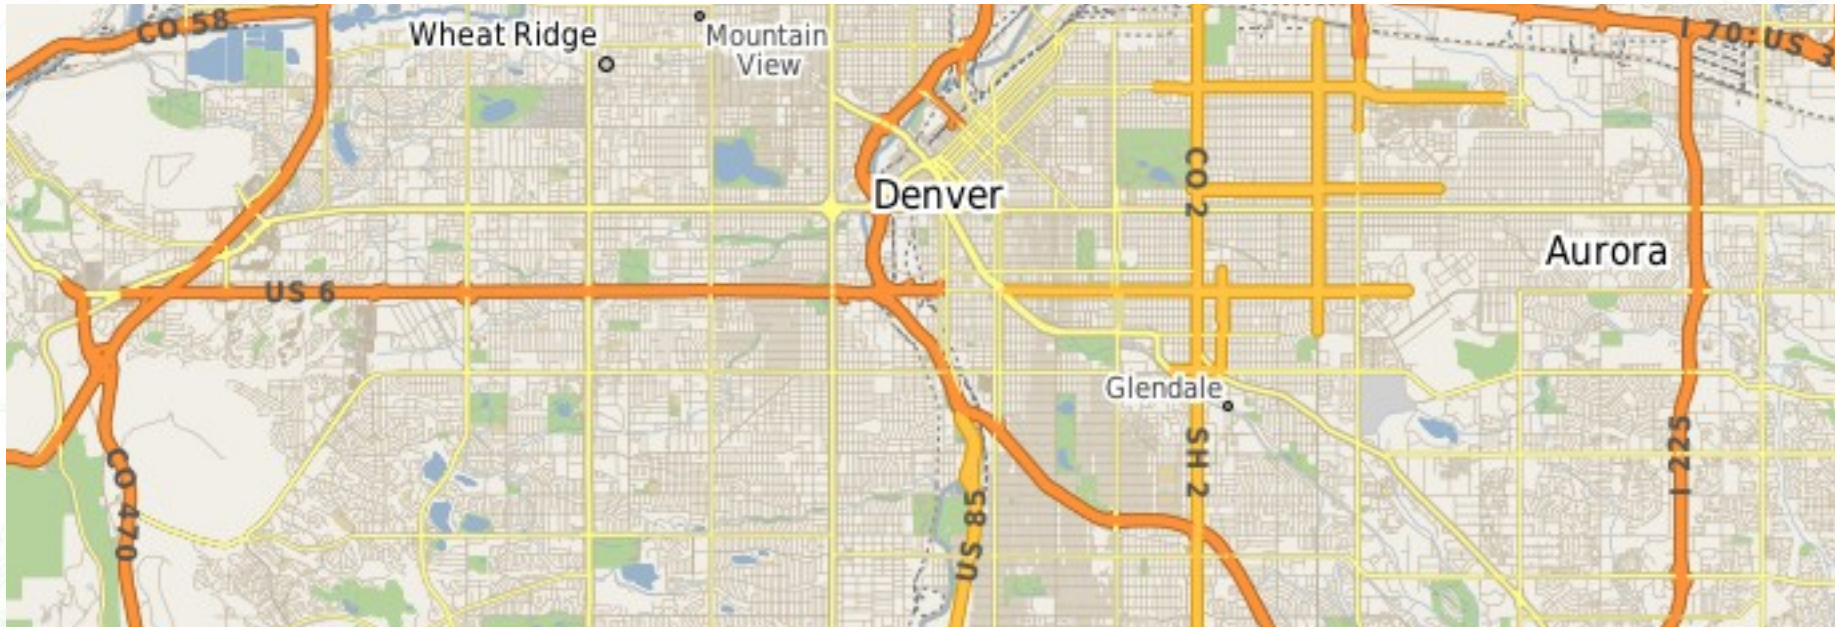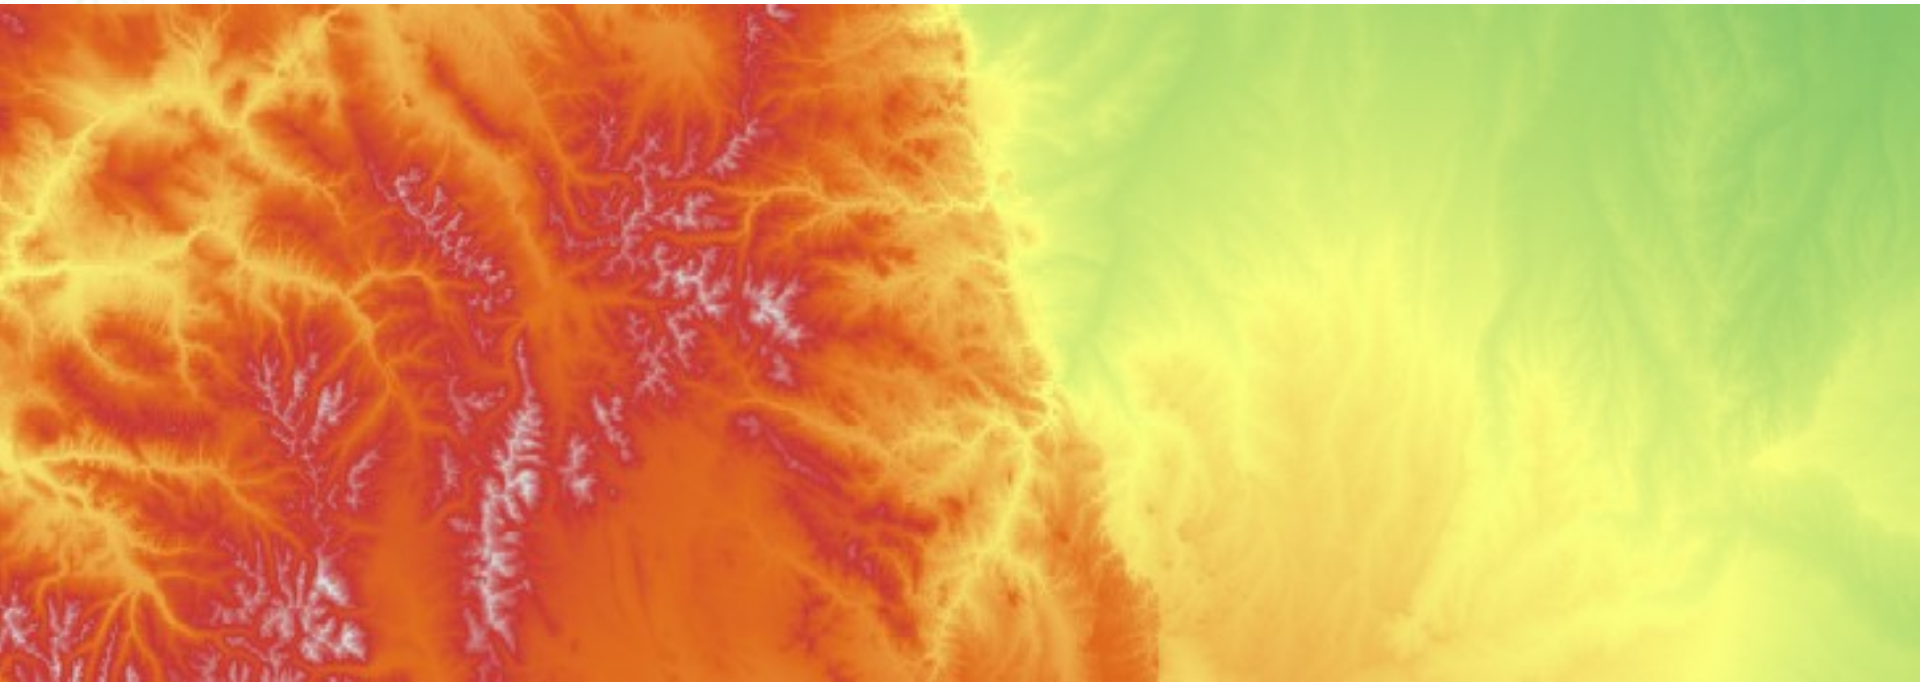
Datasets Used: Extents

Datasets Used: Raster

USGS DEMs NED 1 arc second

- 30m (approx) Resolution
- 16bit Band Interleved Line (BIL)
- Color range dynamically applied based on elevation

Hardware specs

- JMeter:

  Dell Precision Workstation 390 from 9/7/2006
    - Processor, 6300, 1.86, 2M, Core Duo-conroe, Burn 2
    - 2Gb RAM 160 Gb Hard drive 7200 rpm OS: Centos 5.5 i386

- WMS(2):
    - Dell PowerEdge R410   - Ship Date: 7/7/2010
    - Processor: Intel® Xeon® E5630 2.53Ghz, 12M Cache,Turbo, HT, 1066MHz Max Mem
    - 8GB Memory (4x2GB)
    - 2TB 7.2K RPM SATA
    - OS: Windows Server 64bit, Centos 5.5 x86-64

- Database:
    - Gateway E6610D Intel Core2 Duo - E6750 2.66 Ghz
    - 250Gb Hard Drive 7200 rpm, 4Gb Ram
    - OS: Centos 5.5 x86-64

Hardware Configuration

WMS
Linux/Windows

JMeter → Rasters → PostGIS

Bench

Database
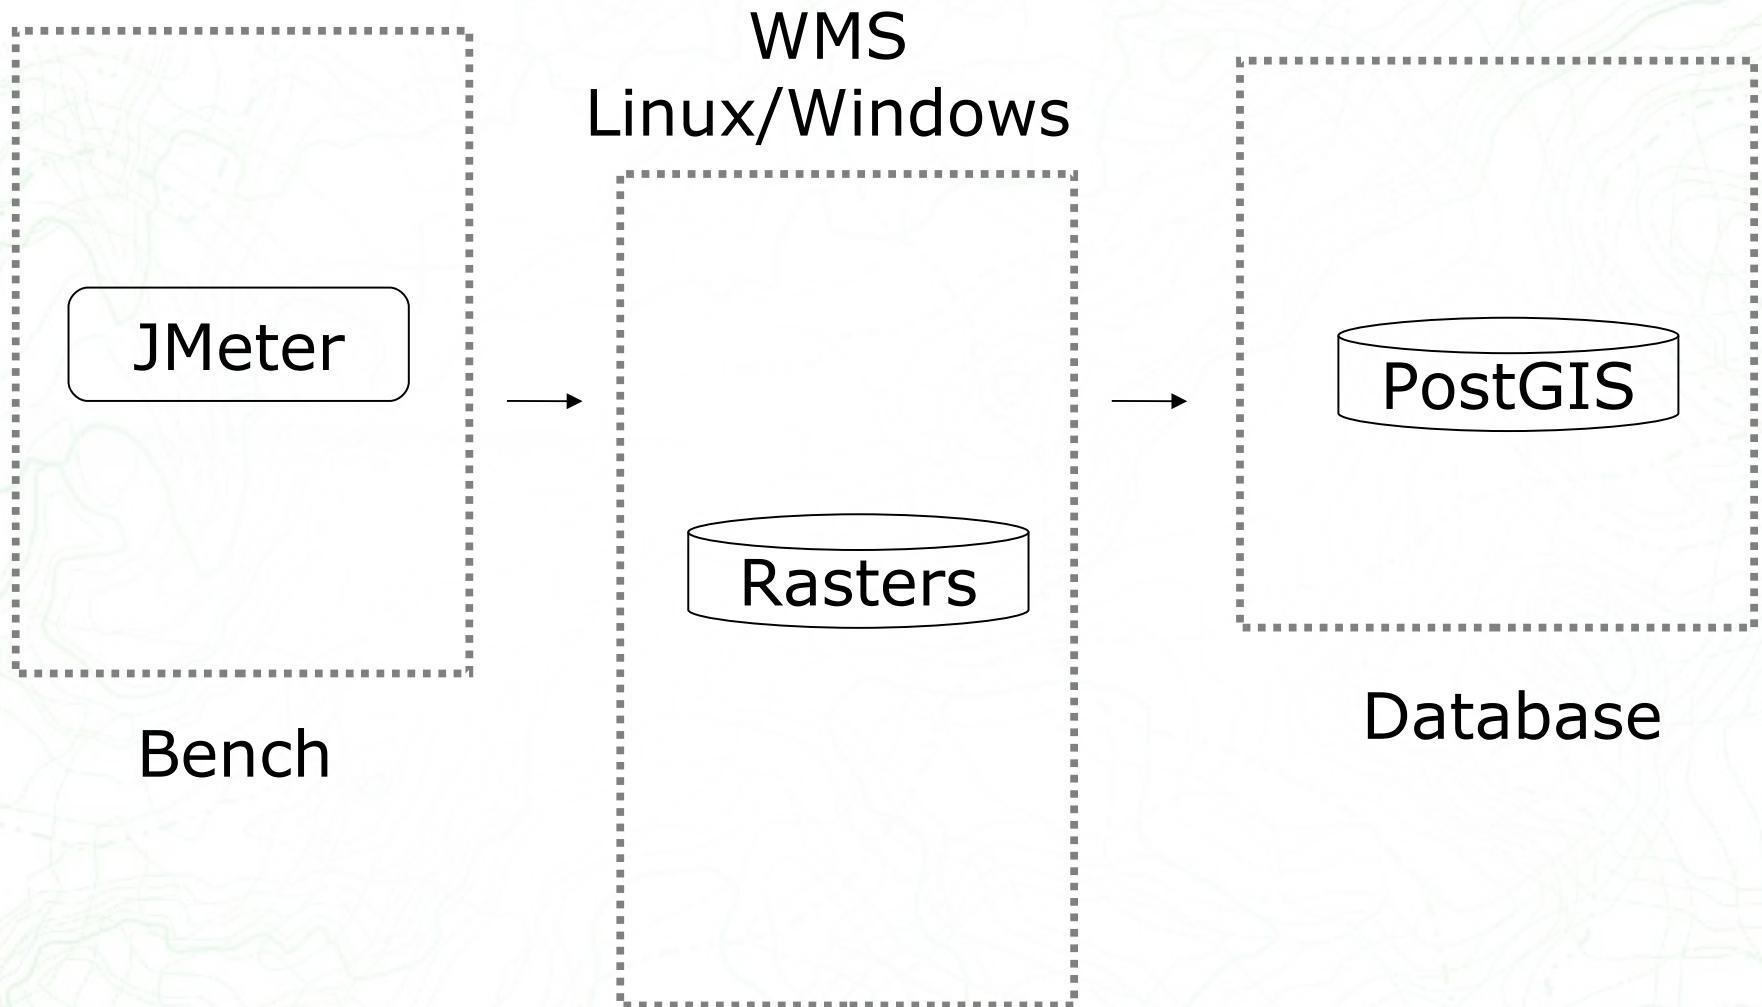
Methodology

- Each test run performs requests with 1, 2, 4, 8, 16, 32, 64, 128 and 256 parallel clients (for a total of 3688 requests)
- Each test uses a random set of requests stored in a CSV file: no two requests in the same run are equal, but all servers perform the same workload
- Two separate tests
  - Normal request
    - The image size (between 64x64 and 1024x768)
    - The geographic envelope (extent)
  - Seed type request
    - Image size fixed (2248x2248)
- Each test is run three times in a row, the results of the third run are used for the comparison: this benchmark assumes full file system caches ("hot" benchmark)
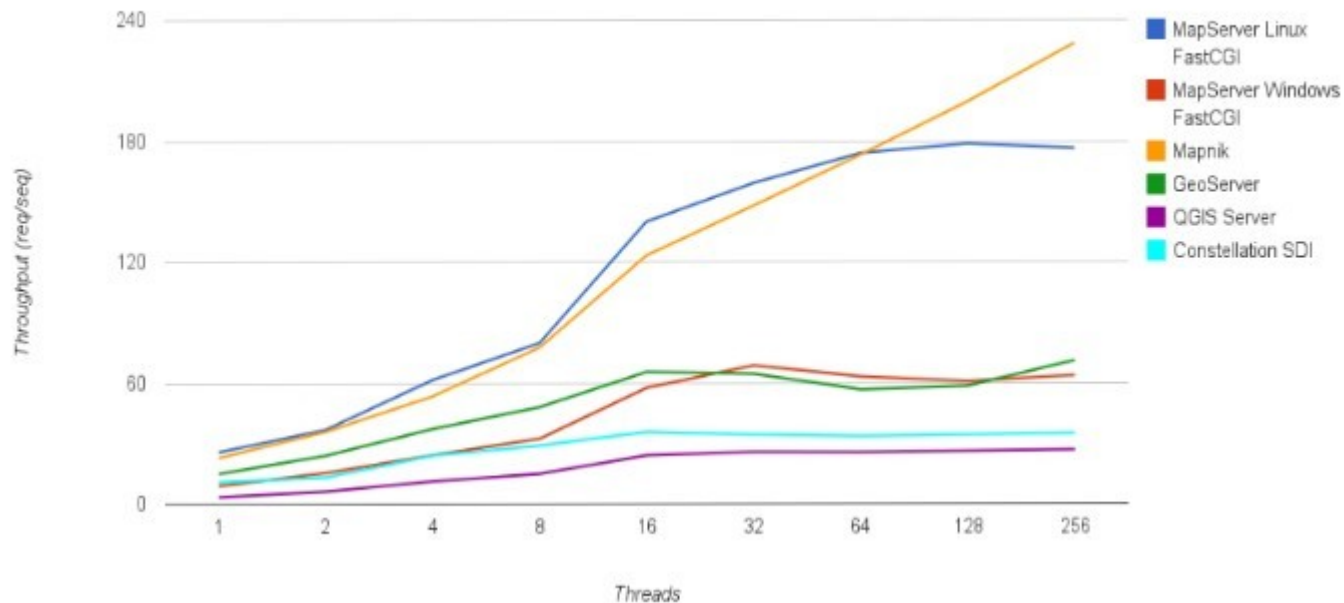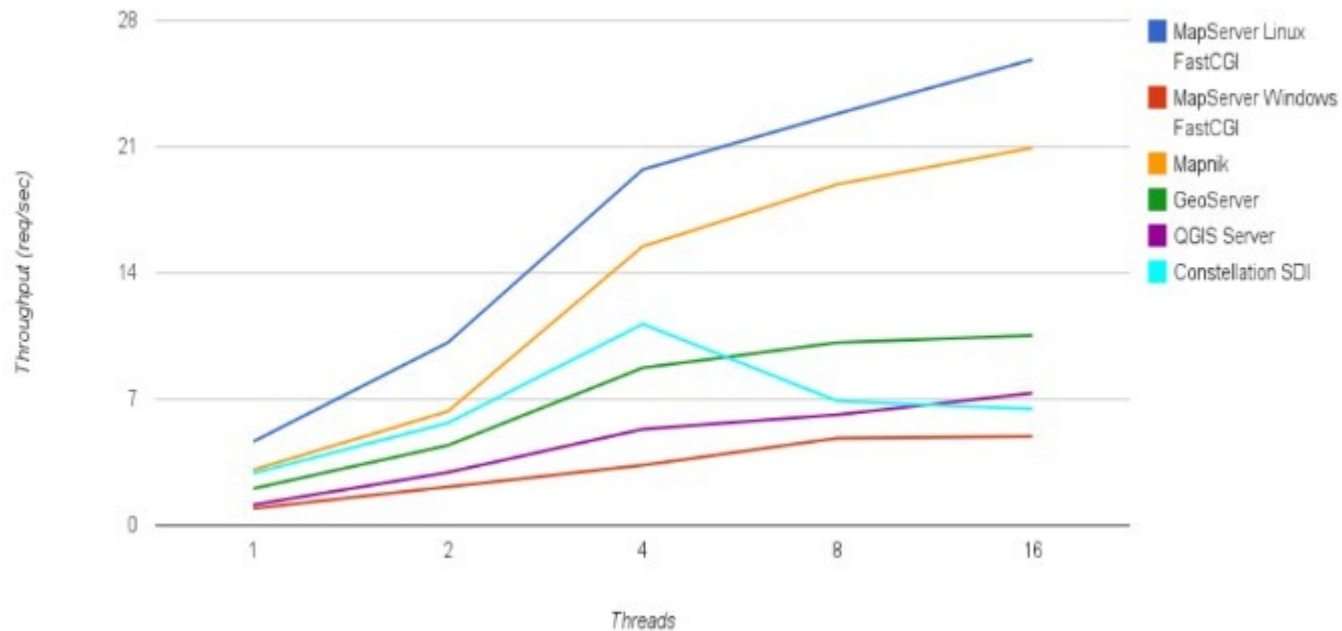
Results

Benchmarking 2011

Vector Results – OSM/PostGIS

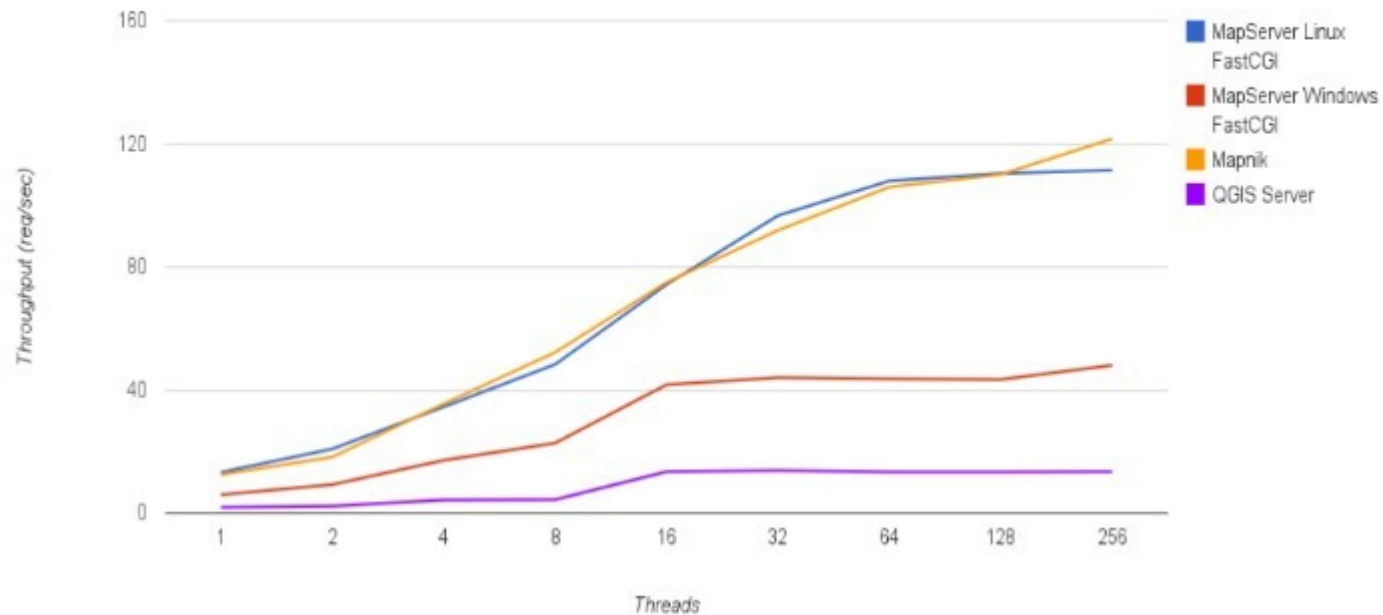# Benchmarking 2011

Vector Seeding Results – OSM/PostGIS

Benchmarking 2011

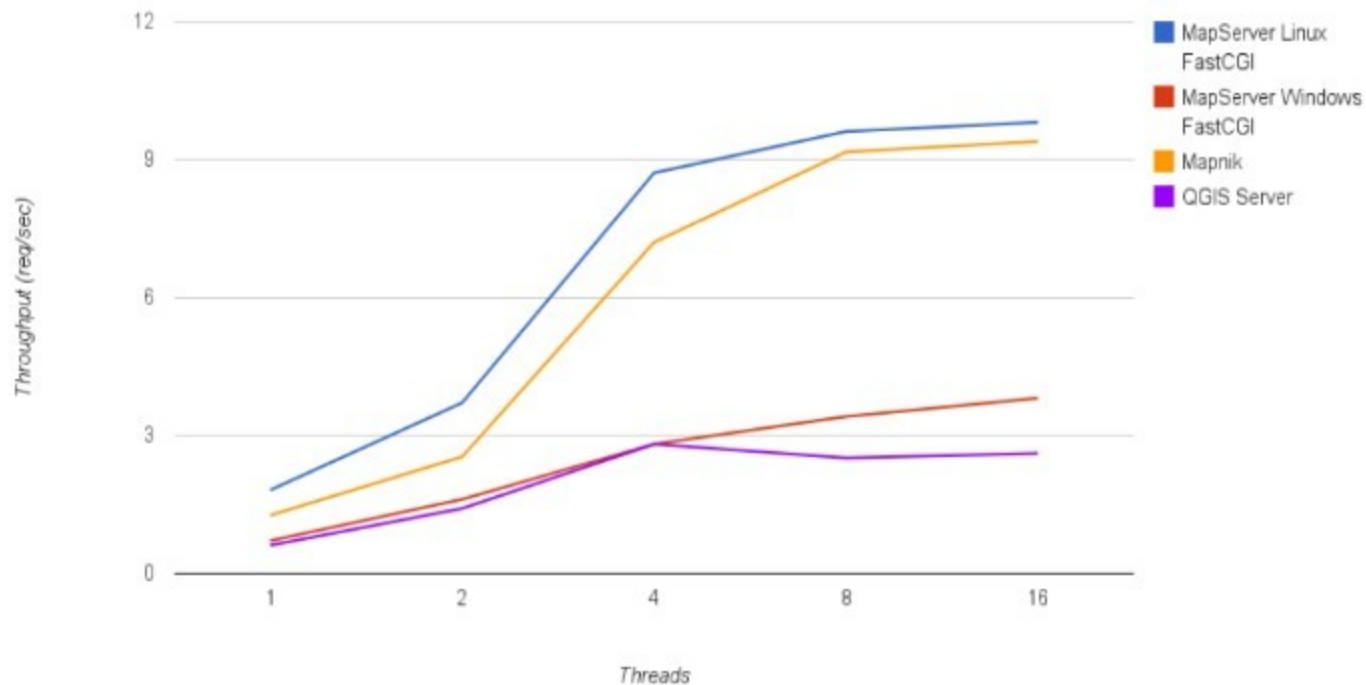# Vector+Raster Results - DEM+OSM

Benchmarking 2011

Vector+Raster Seeding Results - DEM+OSM

Benchmarking 2011
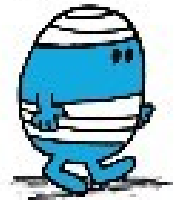
Team Reports

Benchmarking 2011

Cadcorp™

- The Cadcorp GeognoSIS team had to withdraw at the last minute due to a serious family medical emergency

MR. BUMP

- They will resume the tests if at all possible, and publish their results in due course
- So watch this space
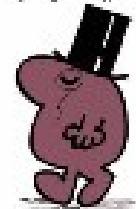- Not that any of *you* care about the results

MR. UPPITY

Benchmarking 2011

- Had to write a MapFile to SLD converter
- Had to write a BIL reader for rasters
    - Work was not completed in time
- Profiling observations
    - CPU usage between 50 and 80%
    - 2/3 of time is spent reading the result of SQL
        - Byte arrays are sent as text (base 64 encoding)
        - Known PostgreSQL-JDBC issue
          (http://postgresql.1045698.n5.nabble.com/bytea-performance-tweak-td4510843.html)

Benchmarking 2011

GeoServer

- Due to lack of time decided to participate only on the vector tests
- No shoot-out specific improvements this year, although didn't want to drop
- Borrowed SLD 1.1 styles from Constellation team (thanks!)
- Hopefully GeoServer keeps doing well
- Same bottleneck: Java2D antialiasing rasterizer. Workaround: load balance an instance each with two cores
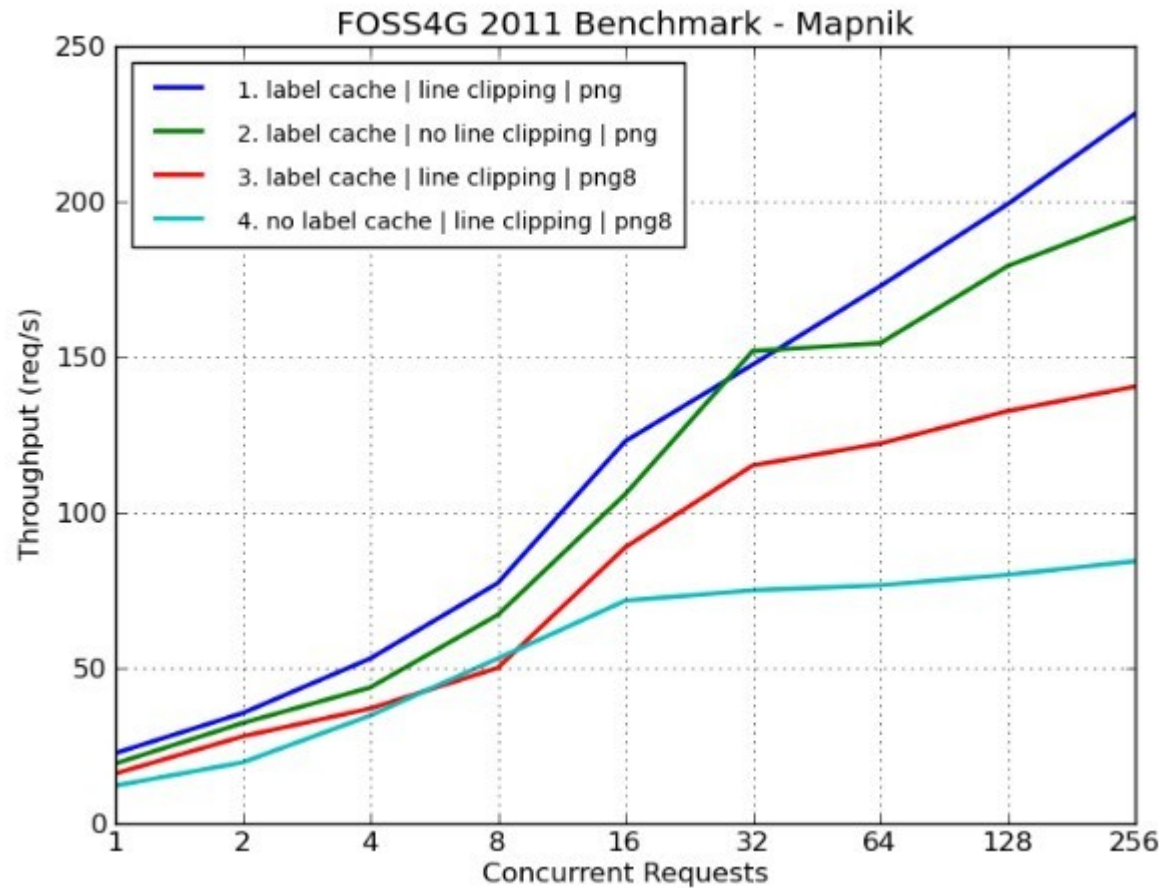
Benchmarking 2011

mapnik

- **Node.js**: New async javascript server: <u>landspeed.js</u>
  - 8 processes (nginx), threadpool of 16 (libeio)
  - >10 r/s than c++ (paleoserver)
- **Labels**:
  - Deferred rendering (avoid extra db queries)
  - Faster placement and halo rendering
- **Raster**: new raster reprojection (using AGG)
- **Encoding**: ability to pass options for zlib/png perf
- **Line drawing**: option for faster rasterization
- **Clipping**: first attempt at polyline clipping
- **Future**: parallel db queries, intelligent feature caching

Thank you: Thomas Bonfort, Konstantin Kaefer, AJ Ashton, Artem Pavlenko, Alberto Valverde, Hermann Kraus, Rob Coup, Simon Tokumine

Benchmarking 2011

FOSS4G 2011 Benchmark - Mapnik

Benchmarking 2011

QGIS Server

- QGS project with 349 layers (200 for labelling)
- Might be the reason for overhead with small tiles
- Improvements of rule-based renderer
- New raster provider with reprojection support performed well
- Data preparation needed, because QGIS requires a unique primary key
- mod_fcgid configuration:
  - Best results with FcgidMaxProcesses = 32
- QGIS uses benchmark server for performance regression testing

Benchmarking 2011 Tweaks & Enhancements

- Png compression parameters
  - Tradeoff between performance and image size
- Apache / mod_fcgid configuration:
  - Run on worker mpm instead of prefork
  - Set FcgidMaxProcessesPerClass to a reasonable value (32) to avoid overwhelming the server with too many processes
    - Default value will eat up postgres connections and lead to failed requests
- Patch to mapserv to not re-parse the (huge!) mapfile at each request
- Patch to MapServer to not apply run-time substitutions (cgi)
- Further optimizations (not done): minimize classification cost
  - Order classes by order of occurence
  - Stop using regexes

Awards

Benchmarking 2011

Last Run Award

GeoServer

Benchmark run at 12:23pm today

Benchmarking 2011

Collaboration Award

Benchmarking 2011

What the shootout is all about

Benchmarking 2011

- Wiki home: http://wiki.osgeo.org/wiki/Benchmarking_2011
- Mailing list: join at http://lists.osgeo.org/mailman/listinfo/benchmarking
- SVN: http://svn.osgeo.org/osgeo/foss4g/benchmarking/wms/2011/

Thank you to all 6 teams!