

# Optimising the Mapnik Rendering Toolchain 2.0



Frederik Ramm  
frederik@remote.org



# Basic Setup

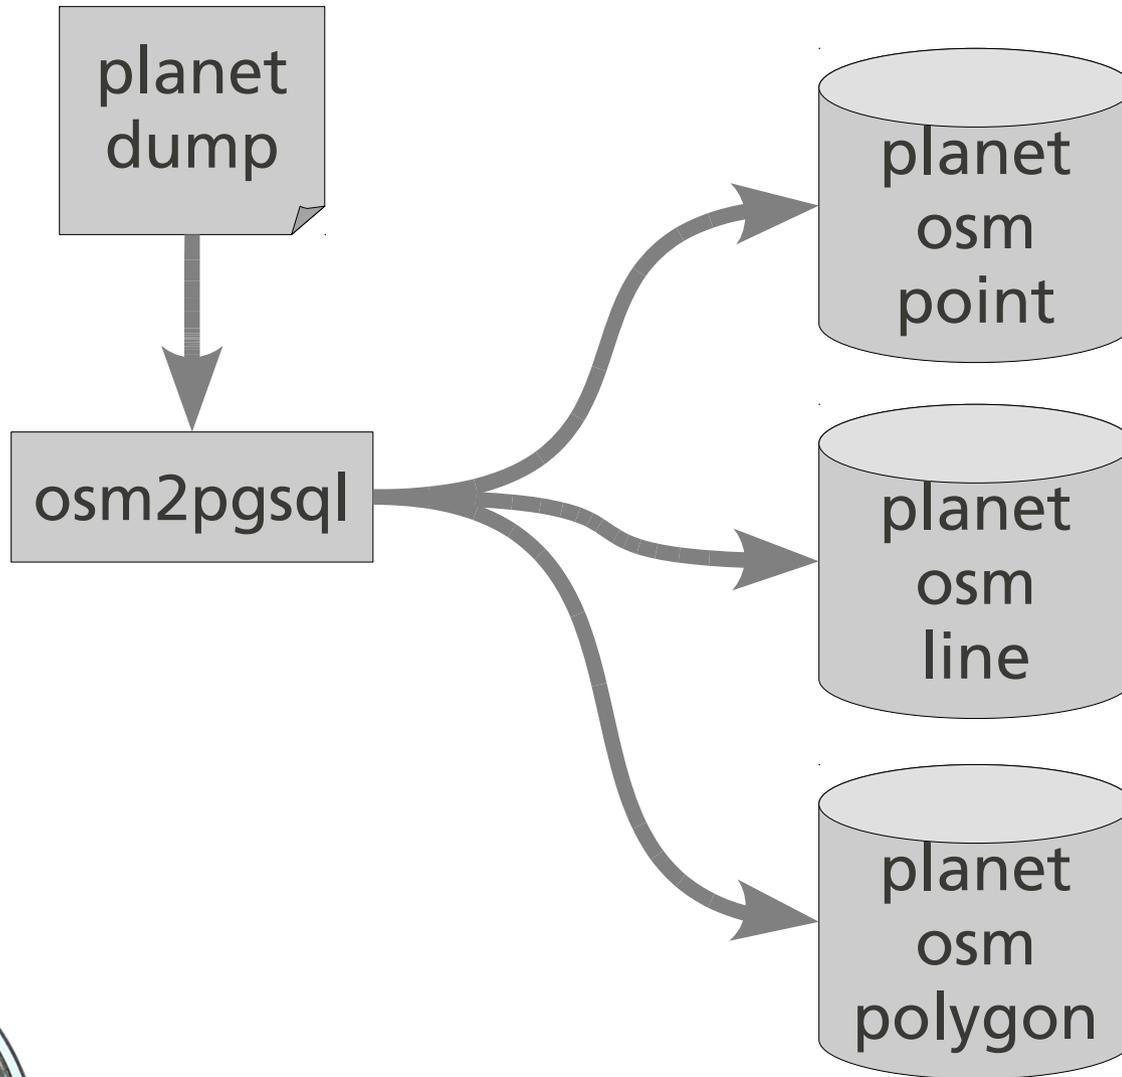
- “Hetzner” dedicated server (US\$ 150/mo)
- Ubuntu Linux
- Mapnik 2.1
- pbf planet file of March 2012
- PostgreSQL 9.1/PostGIS 2.0
- 3x 120 GB SSD
- 32 GB RAM



1.



# What does osm2pgsql do?



high-  
way

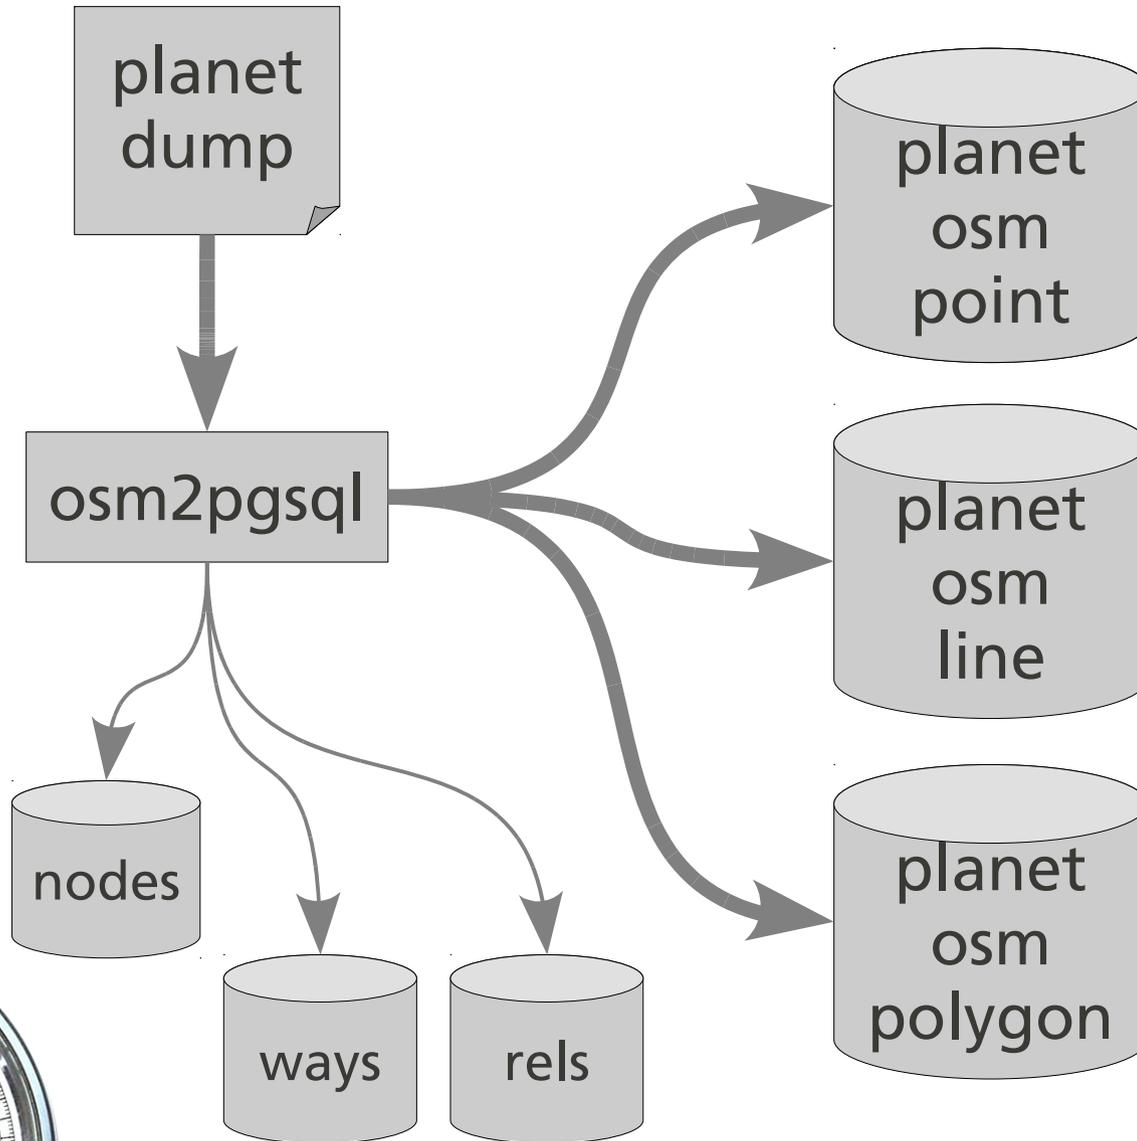
name

one-  
way

...



# osm2pgsql slim mode



highway

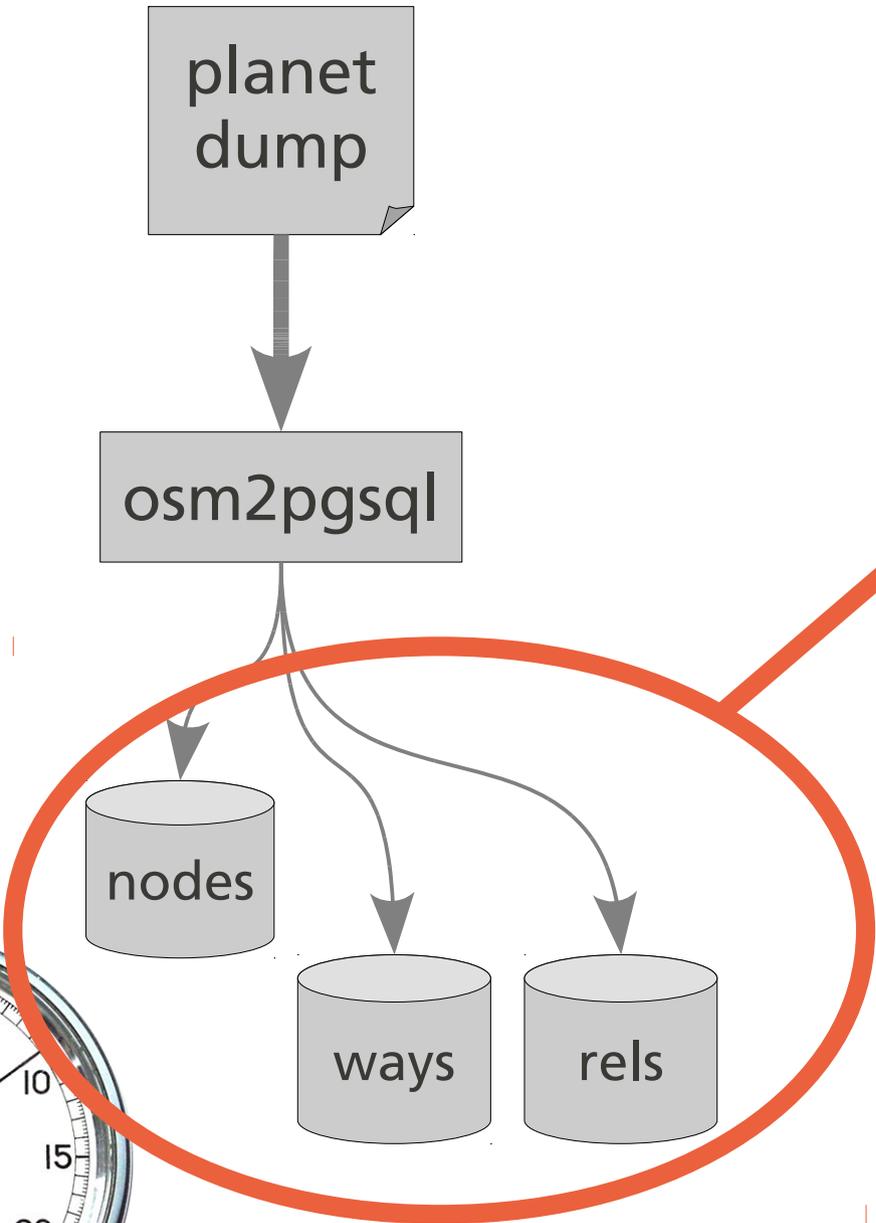
name

one-way

...



# osm2pgsql slim mode



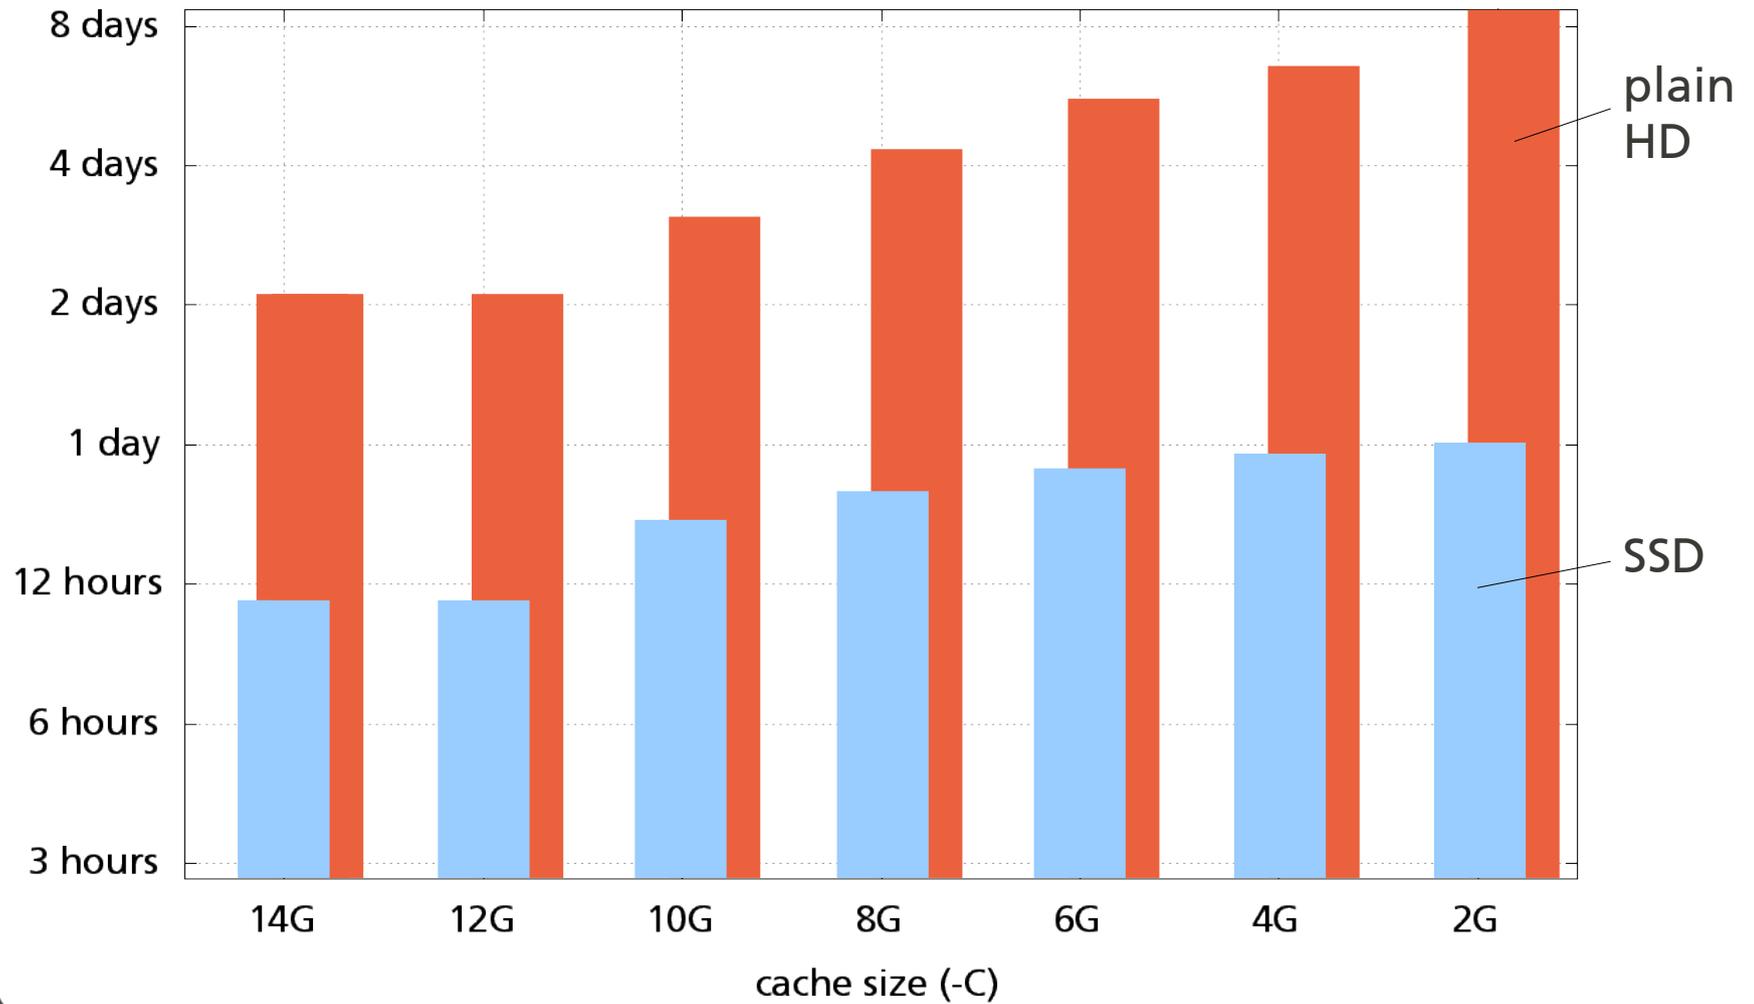
These tables started out as pure "helper" tables for memory-poor systems.

Today they are widely used because they are also a prerequisite for updates!



# Import Time

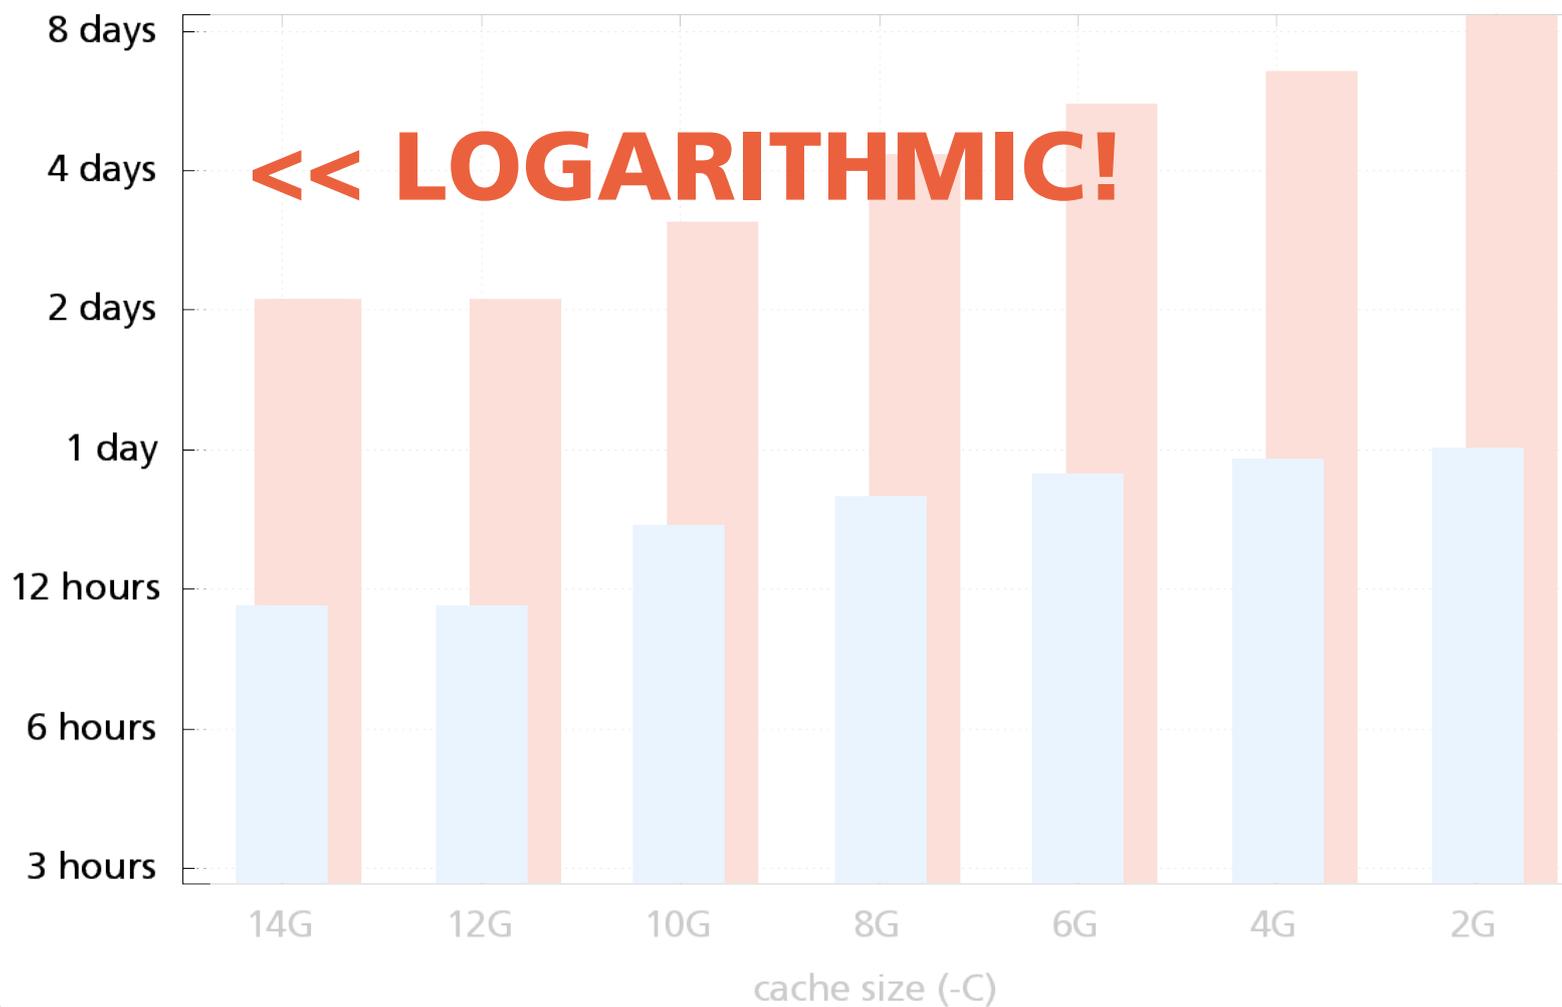
Time for Slim Planet Import



(plain HD: Seagate Barracuda 7200.14 3TB; SSD: 3xOCZ Vertex3)


# Import Time

Time for Slim Planet Import




# Import Time


Time for Slim Planet Import




# So you want to tune your system...



# So you want to tune your system...




# So you want to tune your system...



# Import Time

Time for Slim Planet Import




# So you want to tune your system...



# Import Time

Time for Slim Planet Import



shared_buffers.....	8	1024	4096	8	1024	4096
work_mem.....	1	128	64	1	1	1
maintenance_work_mem.....	4096	256	1024	16	4096	512


PostgreSQL 9.1 memory settings

(all in MB // benchmarks with SSD and -C14G)



# Import Time

Time for Slim Planet Import



synchronous_commits.....	off	off	on	on
fsync.....	off	on	on	on

PostgreSQL 9.1 fsync settings  
(all benchmarks with SSD and -C14G)



# PostgreSQL Tuning for N00bs

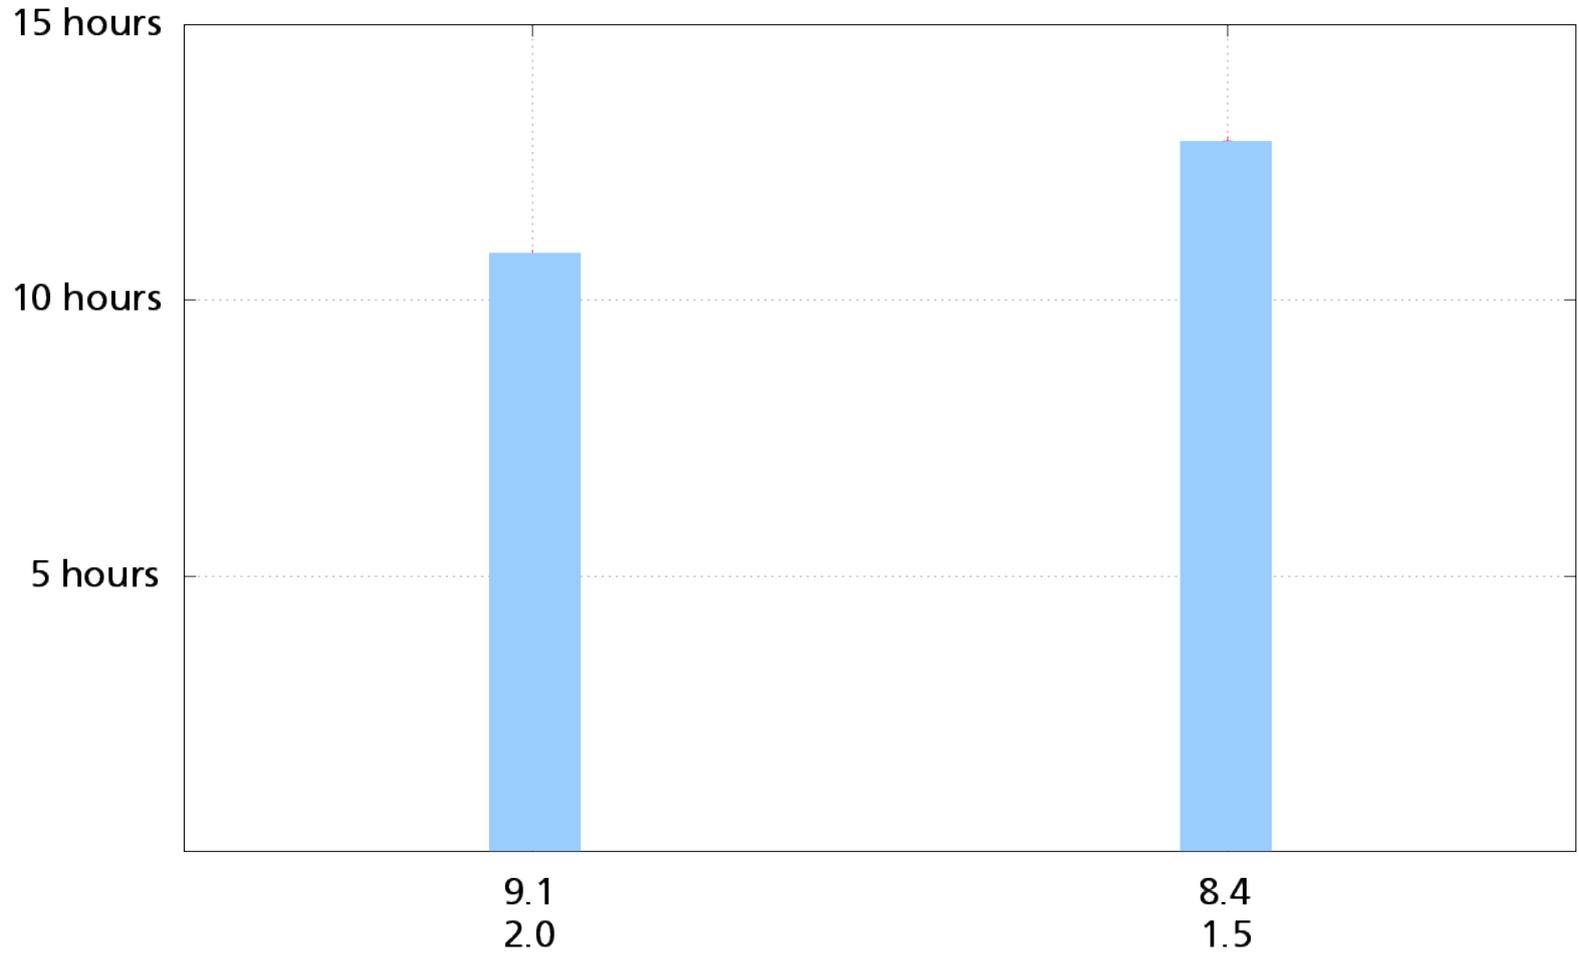
/etc/postgres/9.1/main/postgresql.conf:

option	default	recommended
shared_buffers	24 MB	8 MB
work_mem	1 MB	1 MB
maintenance_work_mem	16 MB	4096 MB
fsync	on	off
autovacuum	on	off (*)
checkpoint_segments		60
random_page_cost	4.0	1.1
effective_cache_size		
effective_io_concurrency	1	



# Import Time

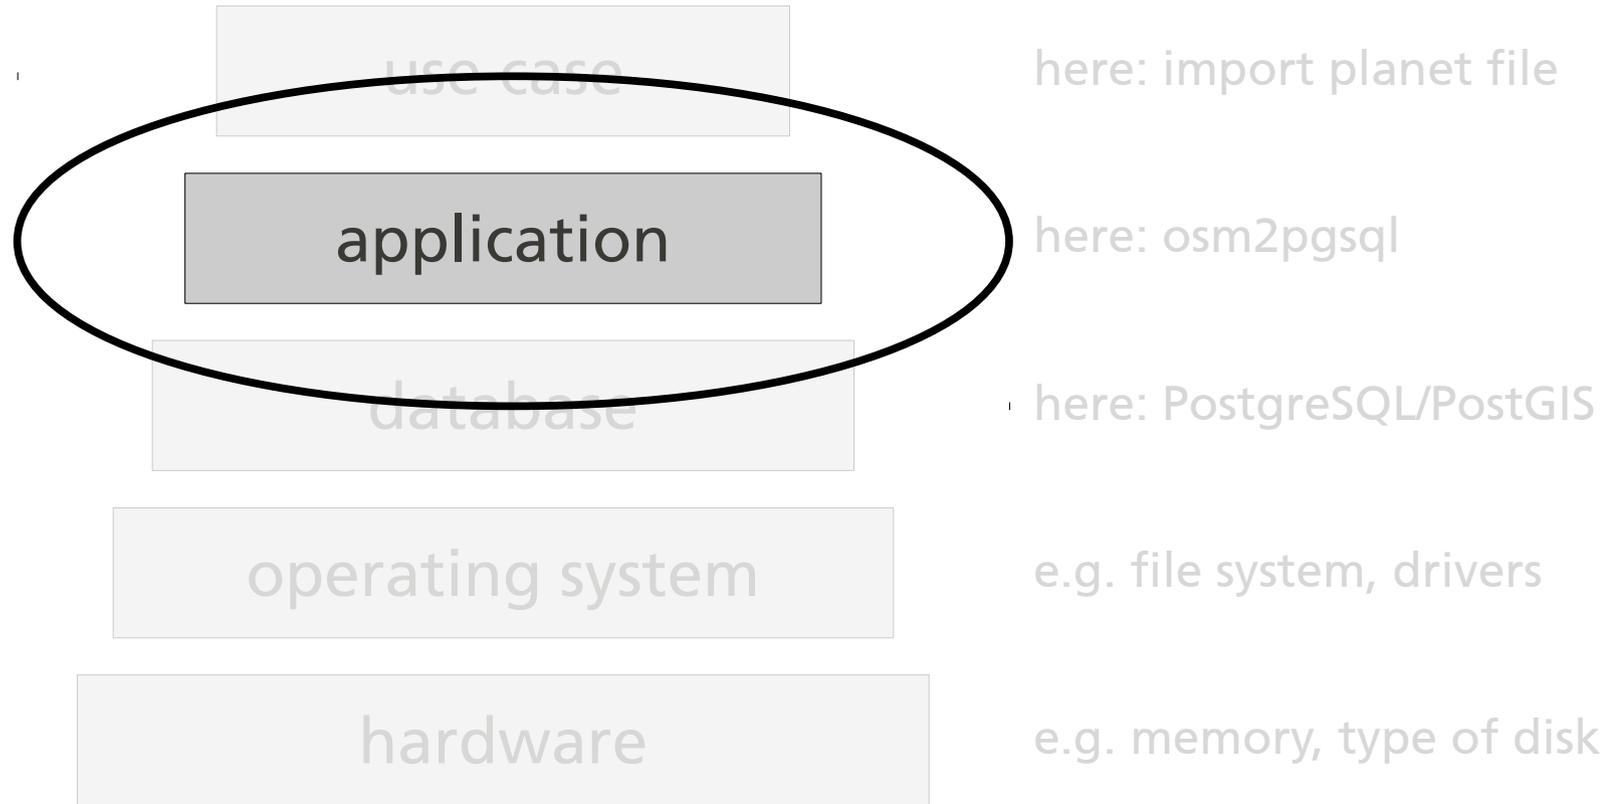
Time for Slim Planet Import



PostgreSQL version, PostGIS version  
(all benchmarks with SSD and -C14G)



# So you want to tune your system...



# Various osm2pgsql options

reference figure: 38956s (short of 11h)

- add reprojection (no -l): + 0.3% (slower)
- use 64bit osm2pgsql: + 11.7% (slower)
- unlogged tables: +17% (slower)
  
- flat node storage: +25% (slower)
- hstore: +10% (slower)



# What is hstore?

- uses a hash column to keep tag key/value combinations instead of tons of columns
- different flavours
- best results: --hstore and --hstore-match-only, with a list of "must-have" tags and a list of "drop" tags
- approx. 10% slower
- disk usage same as before (~ 280G)
- access to all extra tags (e.g. name:xx)



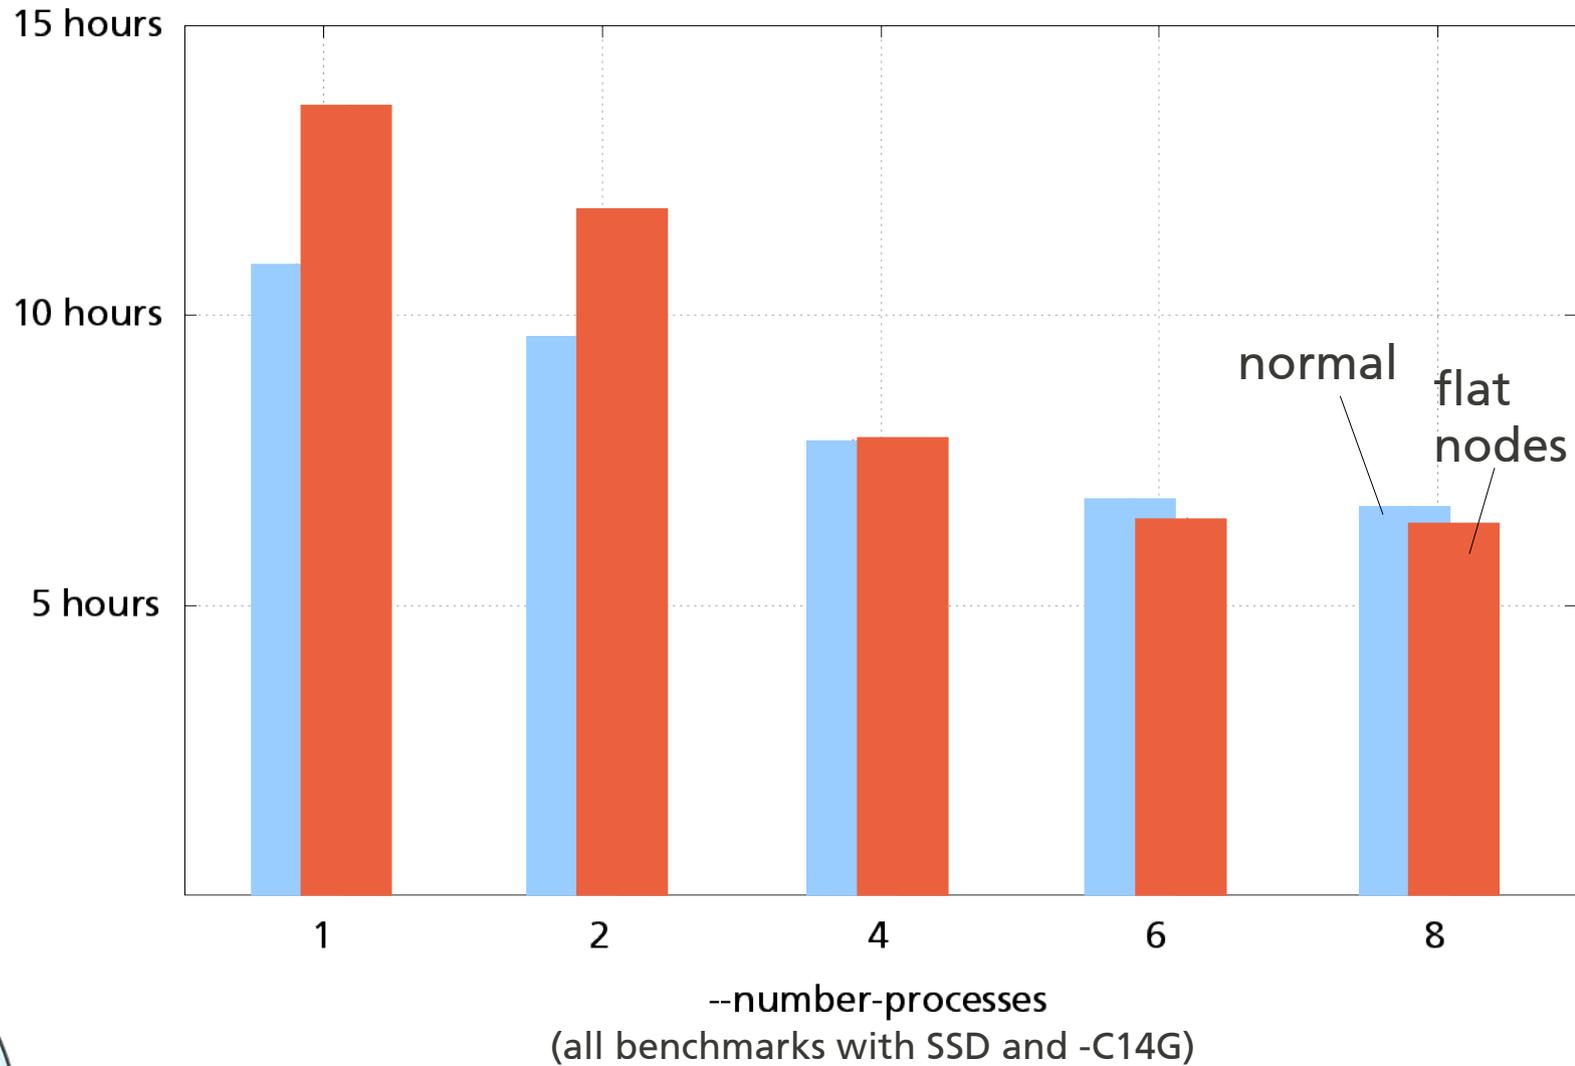
# What is flatnodes?

- instead of storing node lat/lon in a database table, uses a file on disk
- saves about 80G (30%) of disk space
- takes 25% longer but parallelises better
- if you must use a magnetic disk, flatnodes will usually give you a huge performance boost due to much reduced seeking



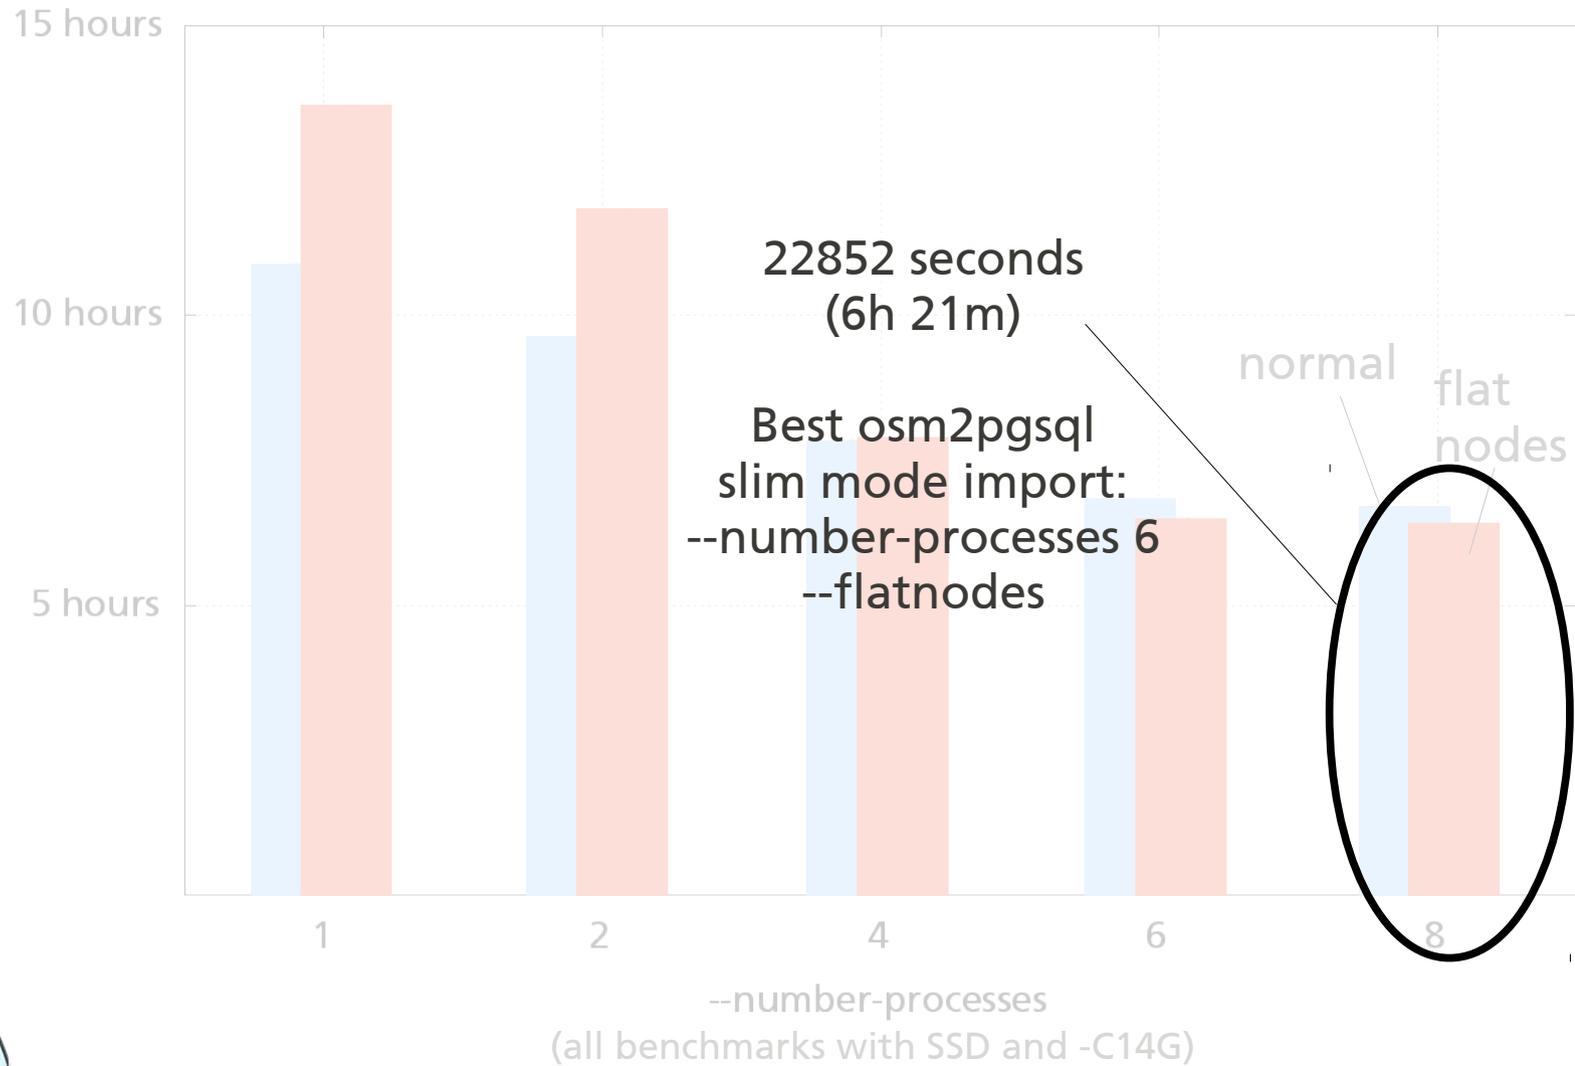
# Parallelisation (8-core machine)

Time for Slim Planet Import




# Parallelisation (8-core machine)

Time for Slim Planet Import



# So you want to tune your system...



# Alternatives to osm2pgsql

- imposm:  
total import time 22893s “out of the box”,  
no updates  
55 GB on disk (+17 temporary)  
(osm2pgsql with `-drop`: 66 GB)  
different table structure;  
50 minutes extra gives you simplified tables
- osmosis:  
unsuitable for rendering
- new OGR driver:  
no benchmarks yet



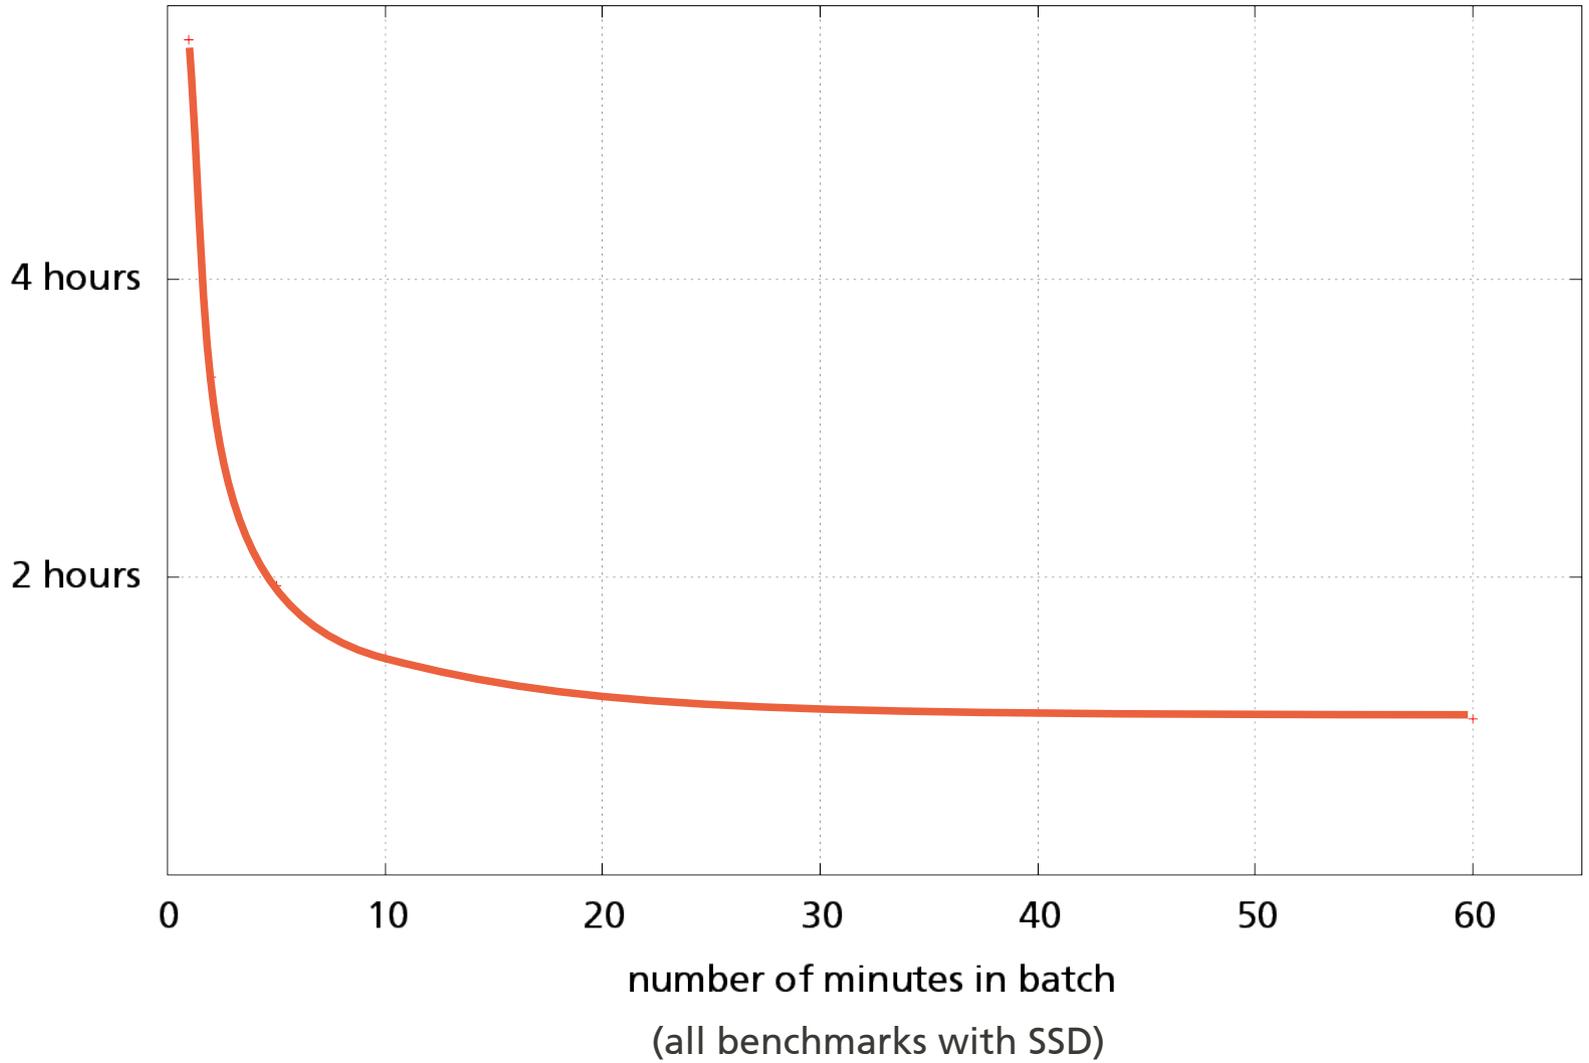
# Slim Mode Updates

- scenario:
  - one day's worth of updates applied to slim import
- basic settings as before



# Update Time

Time for One Day's Worth of Diff Imports



# Slim Mode Updates

- 64 bit: + 10% (slower)
- --number-processes=6: -5% (faster)
- flatnodes: -5% (faster)
- hstore: +/- 0
  
- With HDD, +150% on minutely diffs (spending 14h per day on updates)



# 3.



# Rendering Performance

- scenario:  
a batch of PostgreSQL queries were logged while rendering ~ 17k meta tiles, and replayed in various settings.



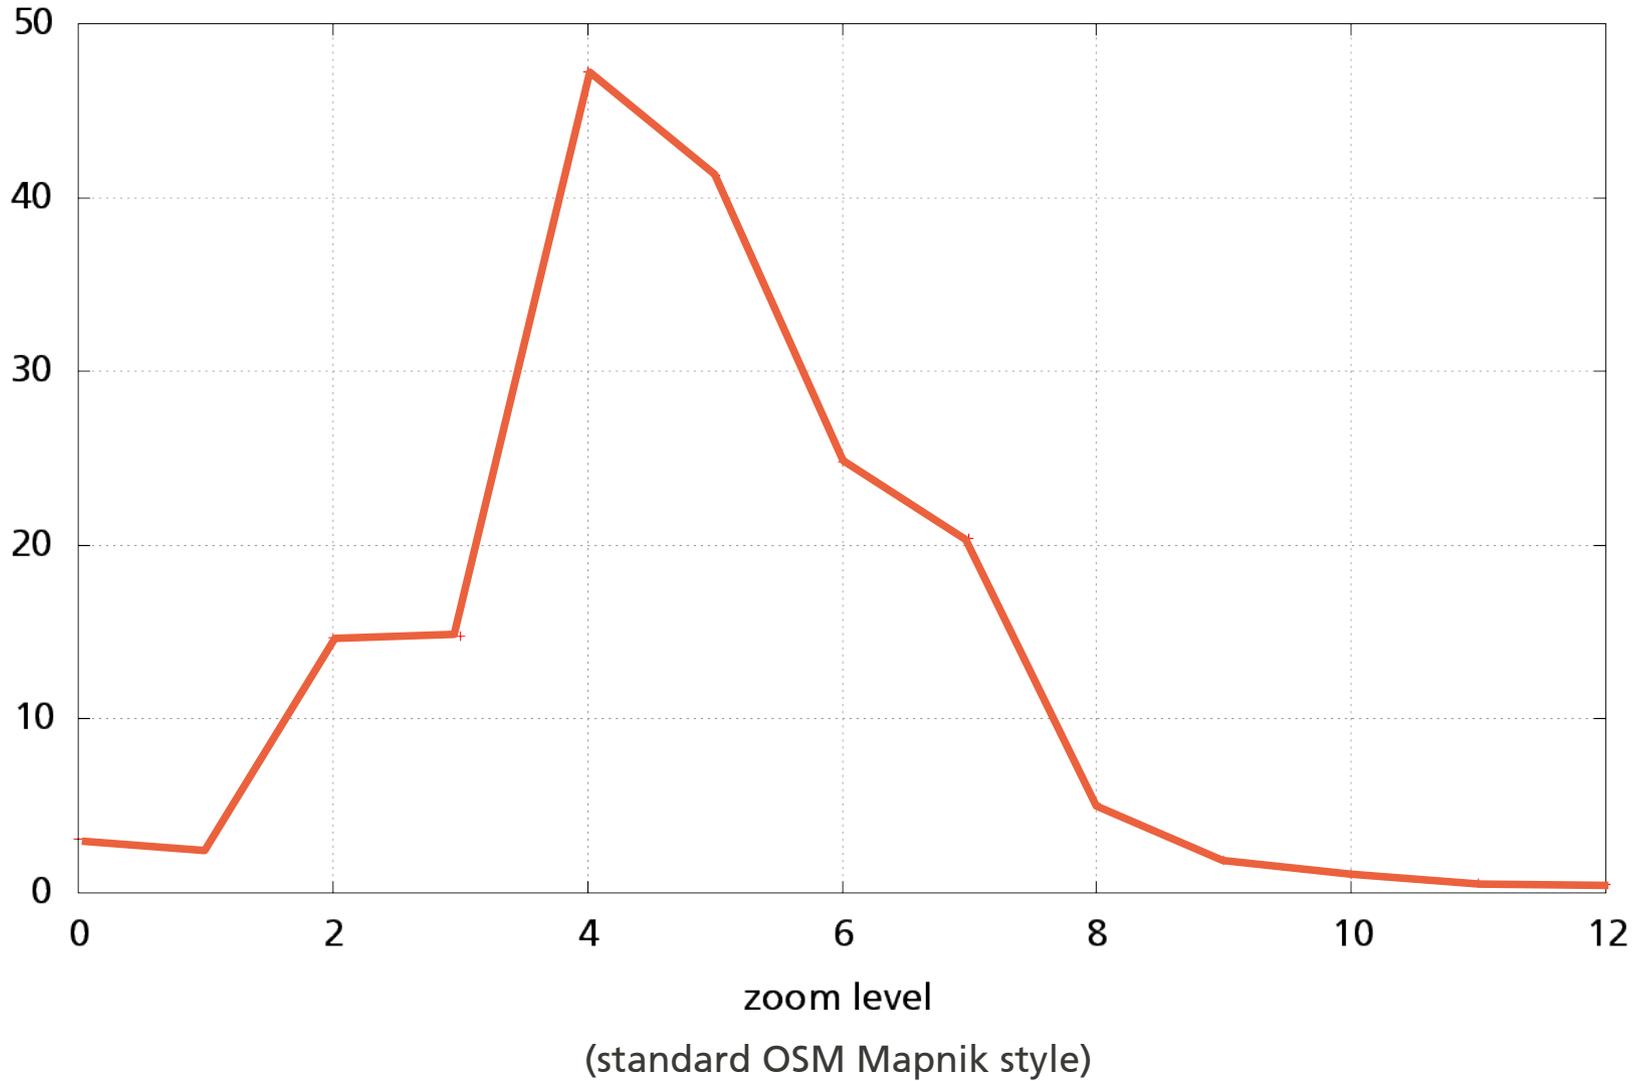
# Rendering Performance

- basic result: 71 minutes (~ 4 MT/s)
- 64bit: + 1%
- HDD: +55%
- flatnode: +/- 0
- hstore and views: + 1%
  
- even deliberately “breaking” PostgreSQL config (different shared\_memory etc) did not change much.



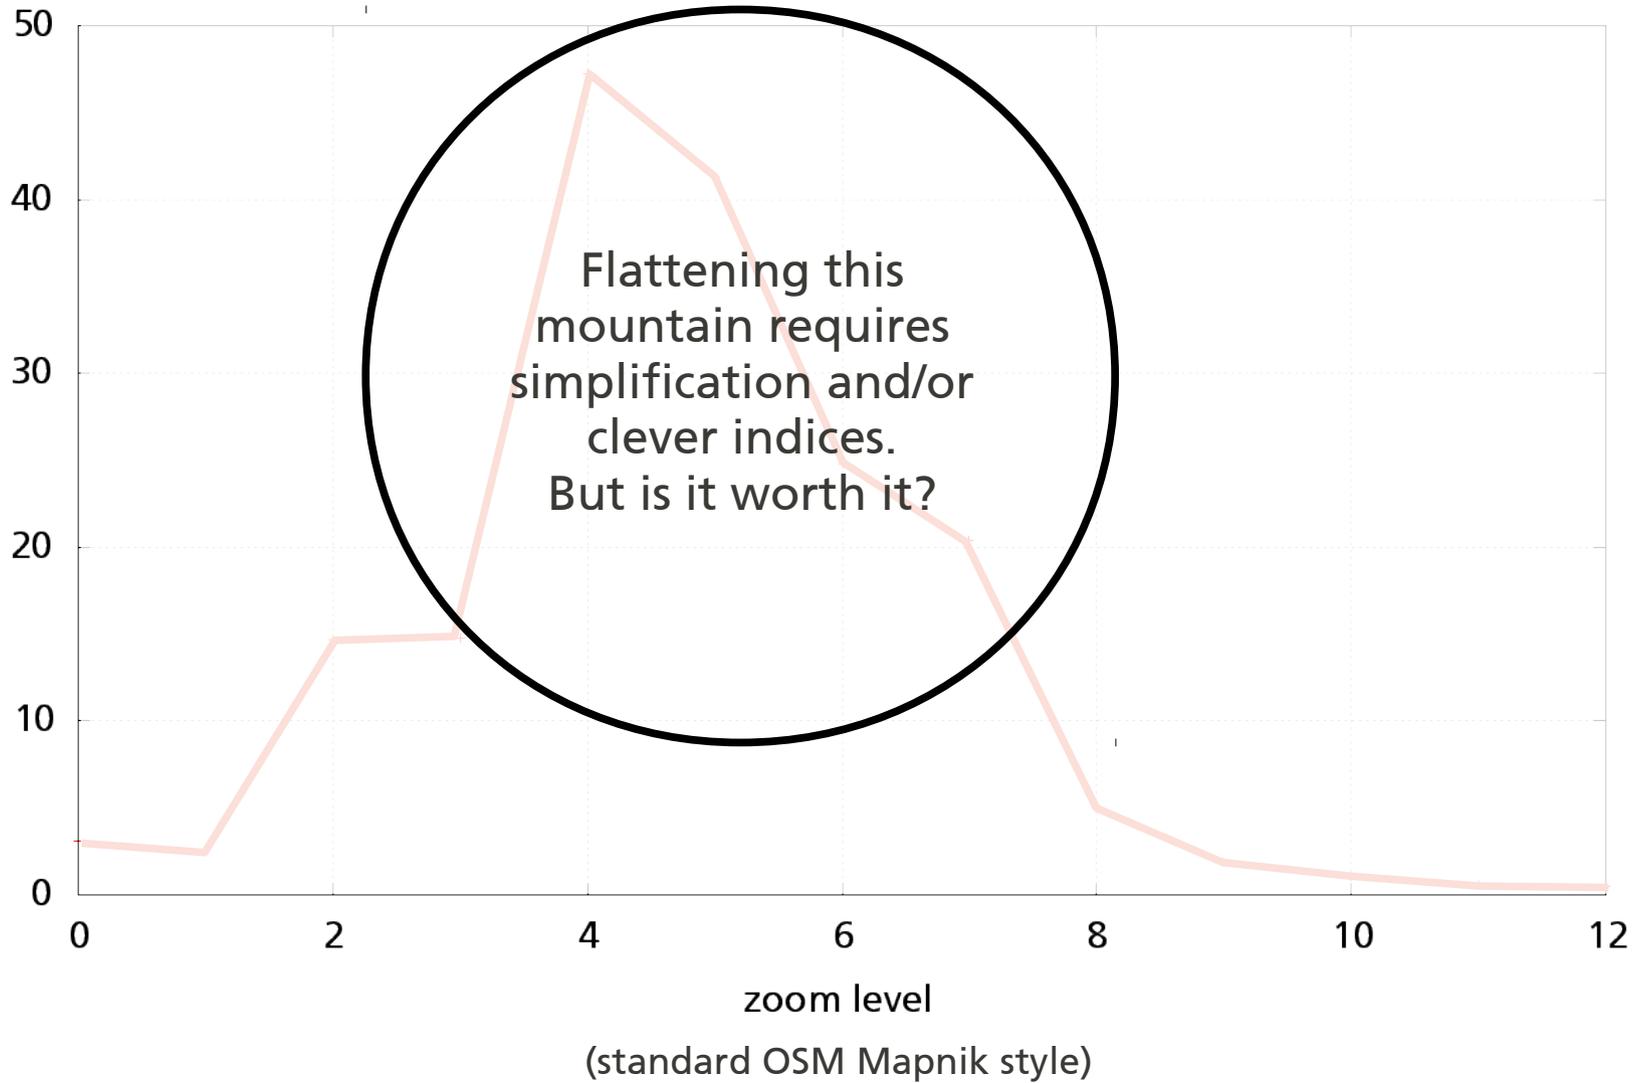
# Rendering Performance

Average Time to Render One Metatile, in seconds



# Rendering Performance

Average Time to Render One Metatile, in seconds



# Rendering Performance

- geometry simplification:  
beforehand (imposm) or on-the-fly (Mapnik option)
- indices:  
use `analyze_postgis_log.pl` from  
[svn.openstreetmap.org/applications/utils/tirex/utils](http://svn.openstreetmap.org/applications/utils/tirex/utils)
- clipping: make sure you don't have giant geometries



# In a Nutshell

- use SSD
- configure PostgreSQL right
- update every 10-15 minutes
- depending on use case, make indexes and simplify geometries



# Thank you

Frederik Ramm  
frederik@remote.org

