

Scaling & PostgreSQL

at myYearbook.com

michael.glaesemann@myyearbook.com

Philly Lambda
2010-07-28

myYearbook.com

casual social network

founded in 2006

Google Analytics

by gender

52% female

48% male

by age

37% 13–17

27% 18–24

13% 25–34

12% 35–44

11% 45+

myYearbook Teens and Twitter Survey (August 2009, ~%)

myYearbook

70% meet people

50% play games

40% keep in touch with friends

35% flirt/date

Facebook

80% keep in touch with friends

35% meet people

30% share photos

25% play games

comScore Teens Category (April 2009)

rank	site	visits (K)	uniques (K)	minutes (M)	page views (M)
1	myYearbook	55,808	4,604	851	1,630
2	MEEZ	8,629	1,407	226	469
3	Zwinky	11,558	3,691	153	108
4	Hearst Teen Network	6,940	2,314	53	117
5	Quizilla	7,924	2,058	75	75

comScore page views
July 2009

rank	site	views (M)
20	GaiaOnline	1,105
21	Chase	1,056
22	ESPN	984
23	myYearbook	953
24	Wikipedia	903
25	Onemanga	840
26	Mapquest	833
27	Foxsports	815

comScore time spent
July 2009

rank	site	minutes (M)
20	Amazon	806
21	CNN	744
22	GaiaOnline	709
23	Bing	707
24	MSNBC	691
25	myYearbook	678
26	Iwin	670
27	NickJr	665

PostgreSQL

UTF-8

correct,
then fast

Codd

relational theory
values & operators

normalization

logical

I do declare!

SQL

Quel

Prolog

Datalog

regular expressions

BOOM

Hive

logical

versus

physical

planner

compiler

thought

word

deed

EXPLAIN

Stonebraker

Ingres

Postgres, Illustra

c-store, Vertica

h-store, VoltDB

Aster Data
Greenplum
EnterpriseDB
Truviso

Big in Japan

Fujitsu

NTT Data

NTT OSS

SRA

BSD

object relational

MVCC


VACUUM

time travel

transactional
DDL

but does it
scale?

fsync



more activity?

views/month 2007 100M → 2009 1.5G

more TPS

more servers

more connections

more configuration

more pain!


Skype

PL/Proxy


pgbouncer


less configuration


fewer connections




fewer connections!

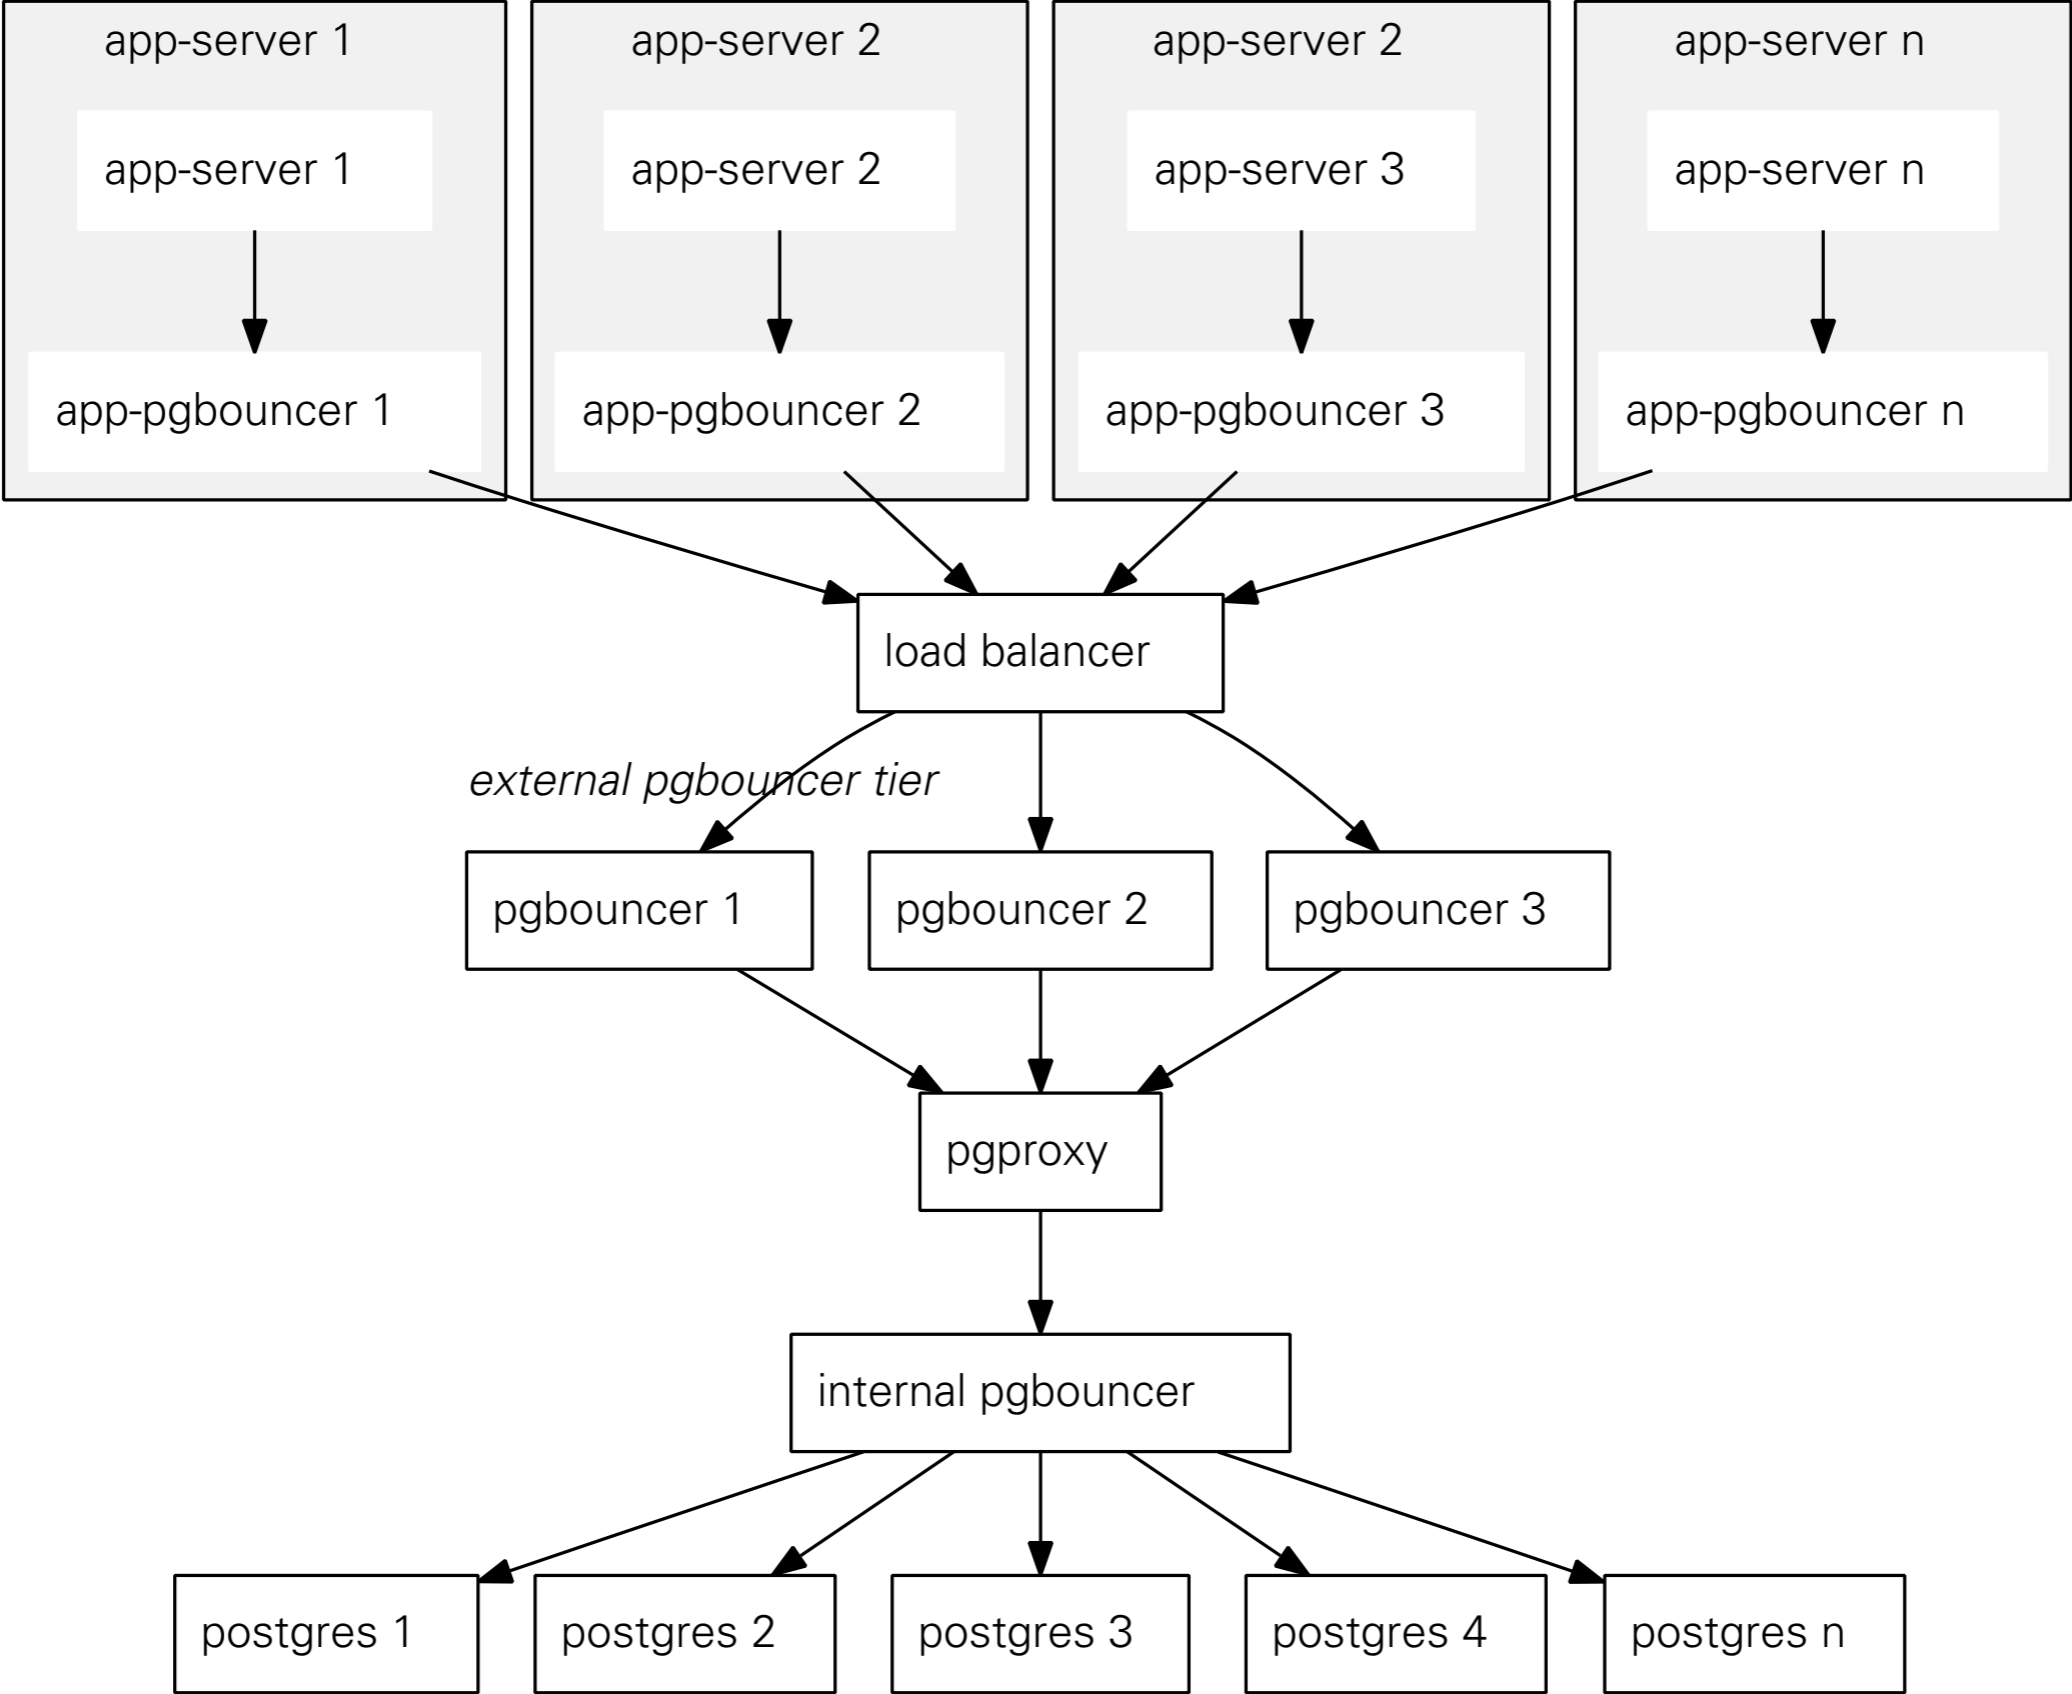


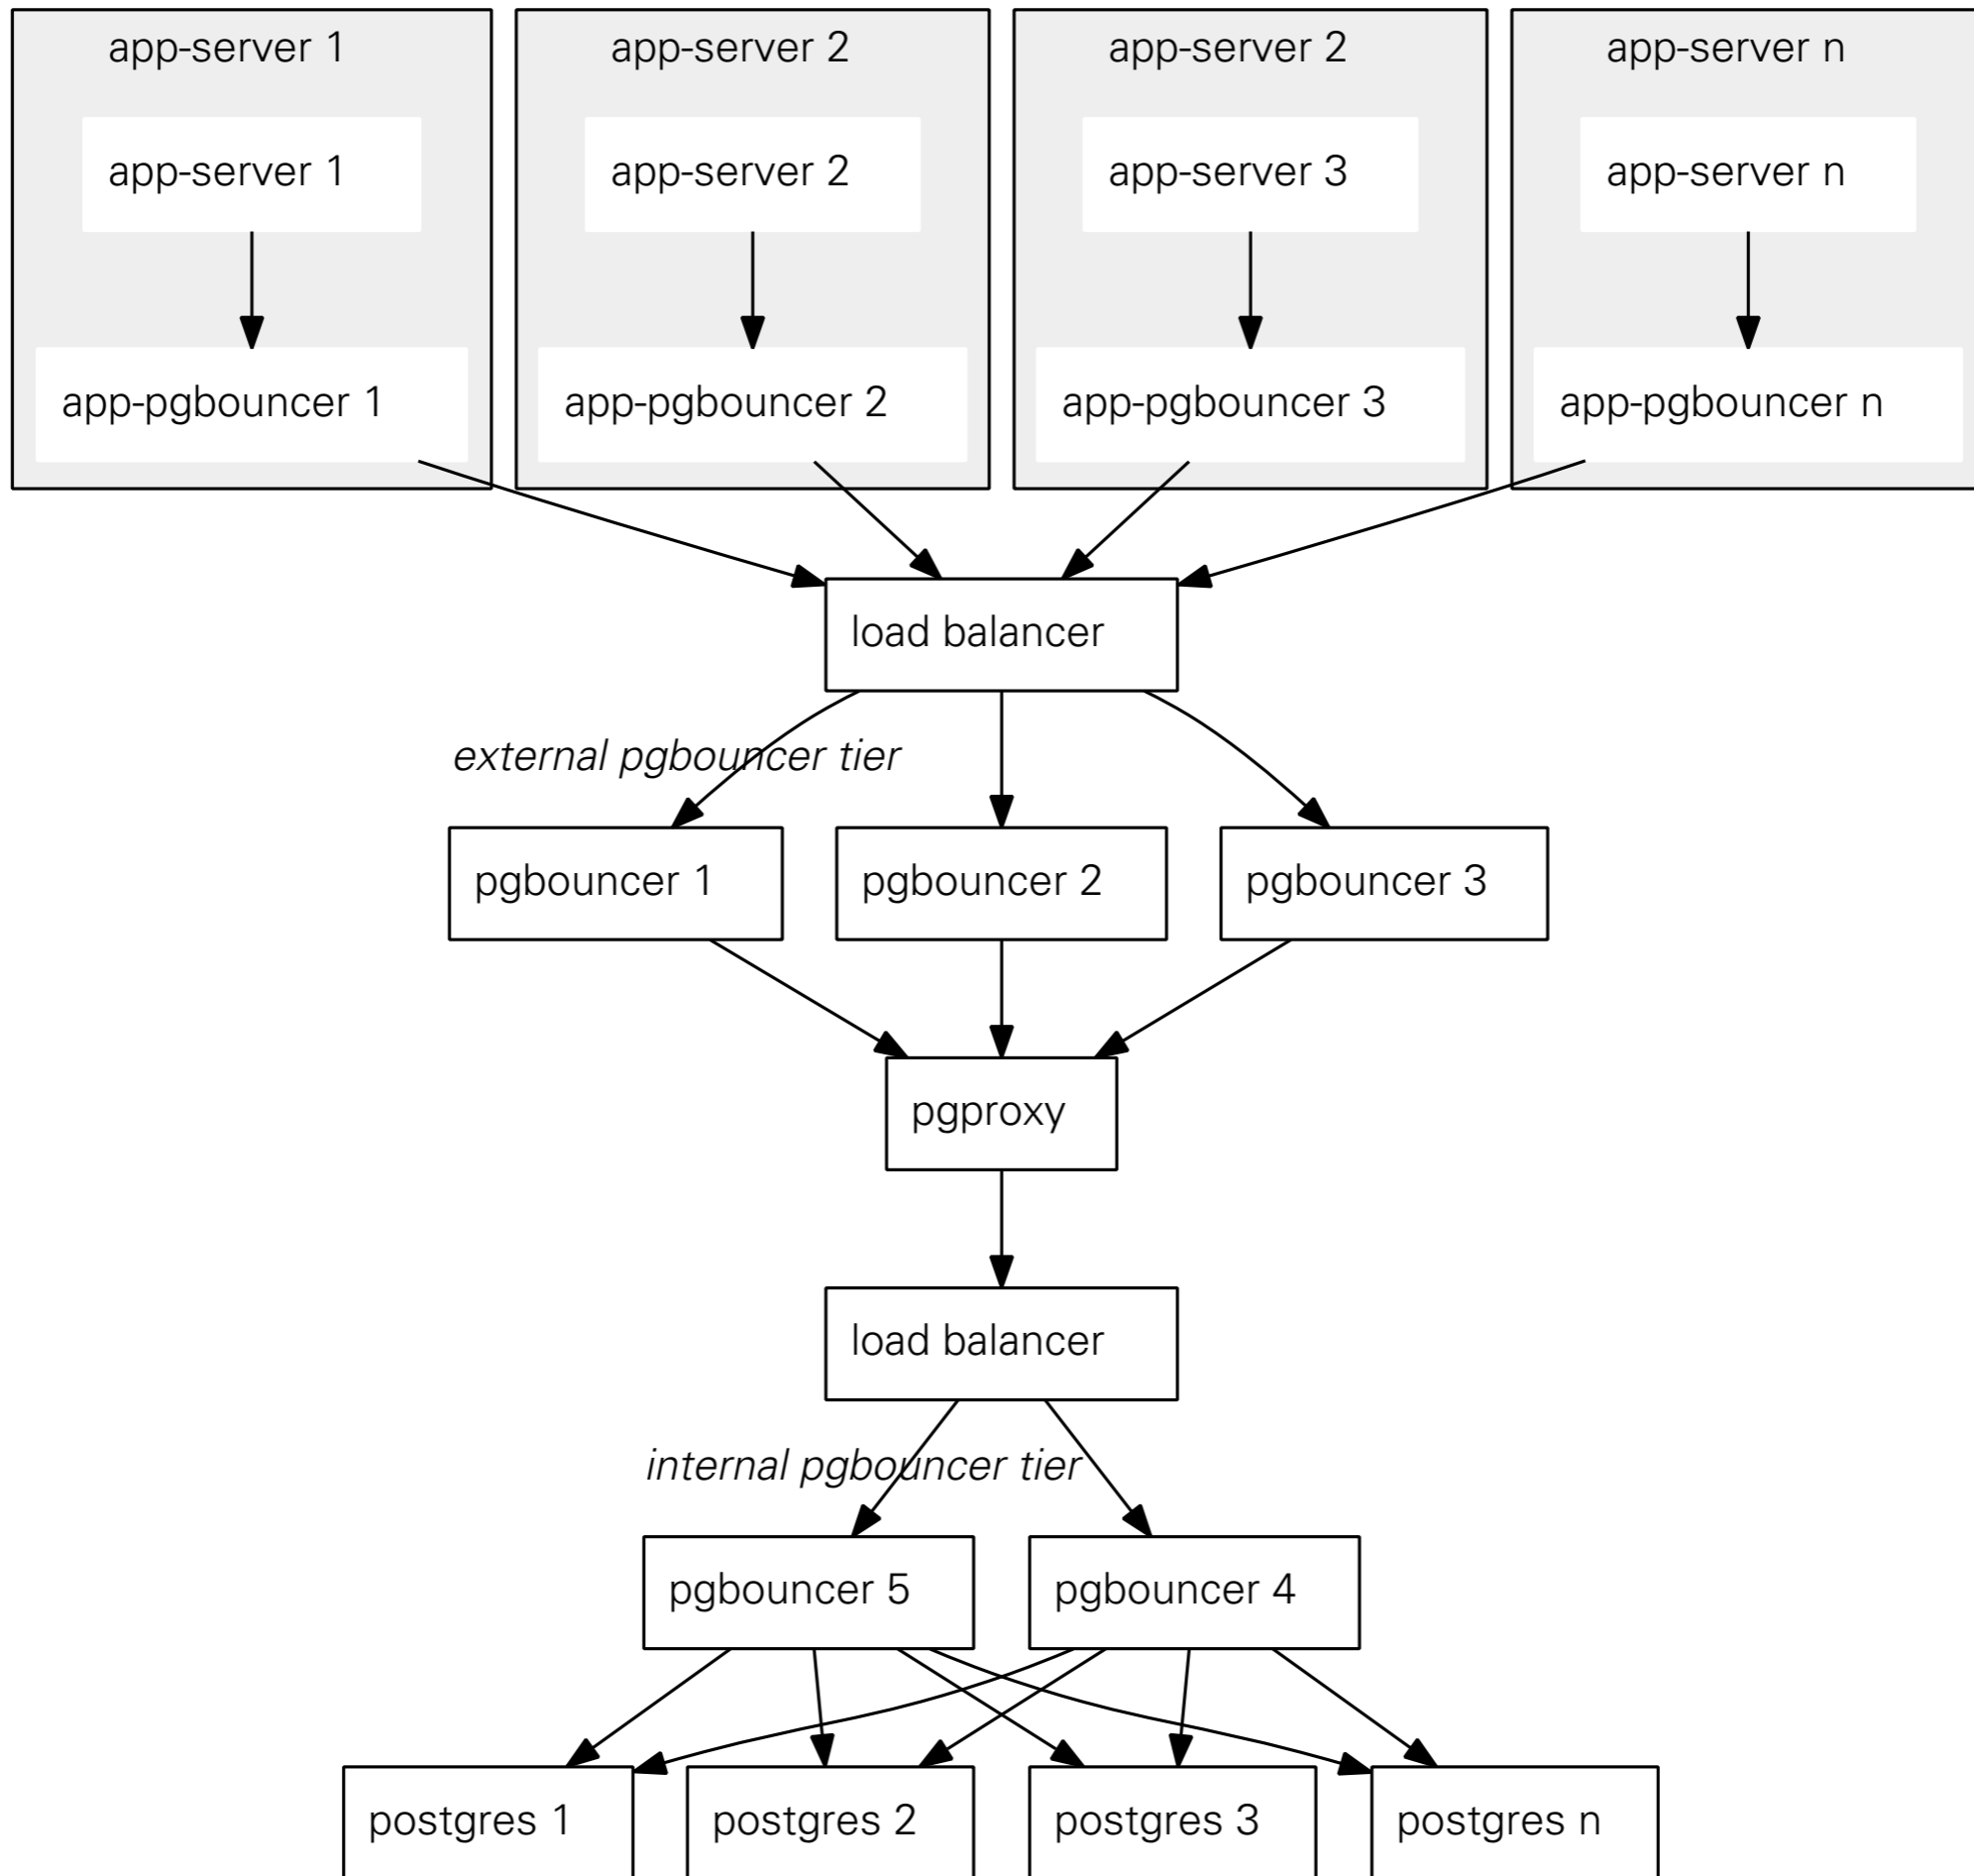


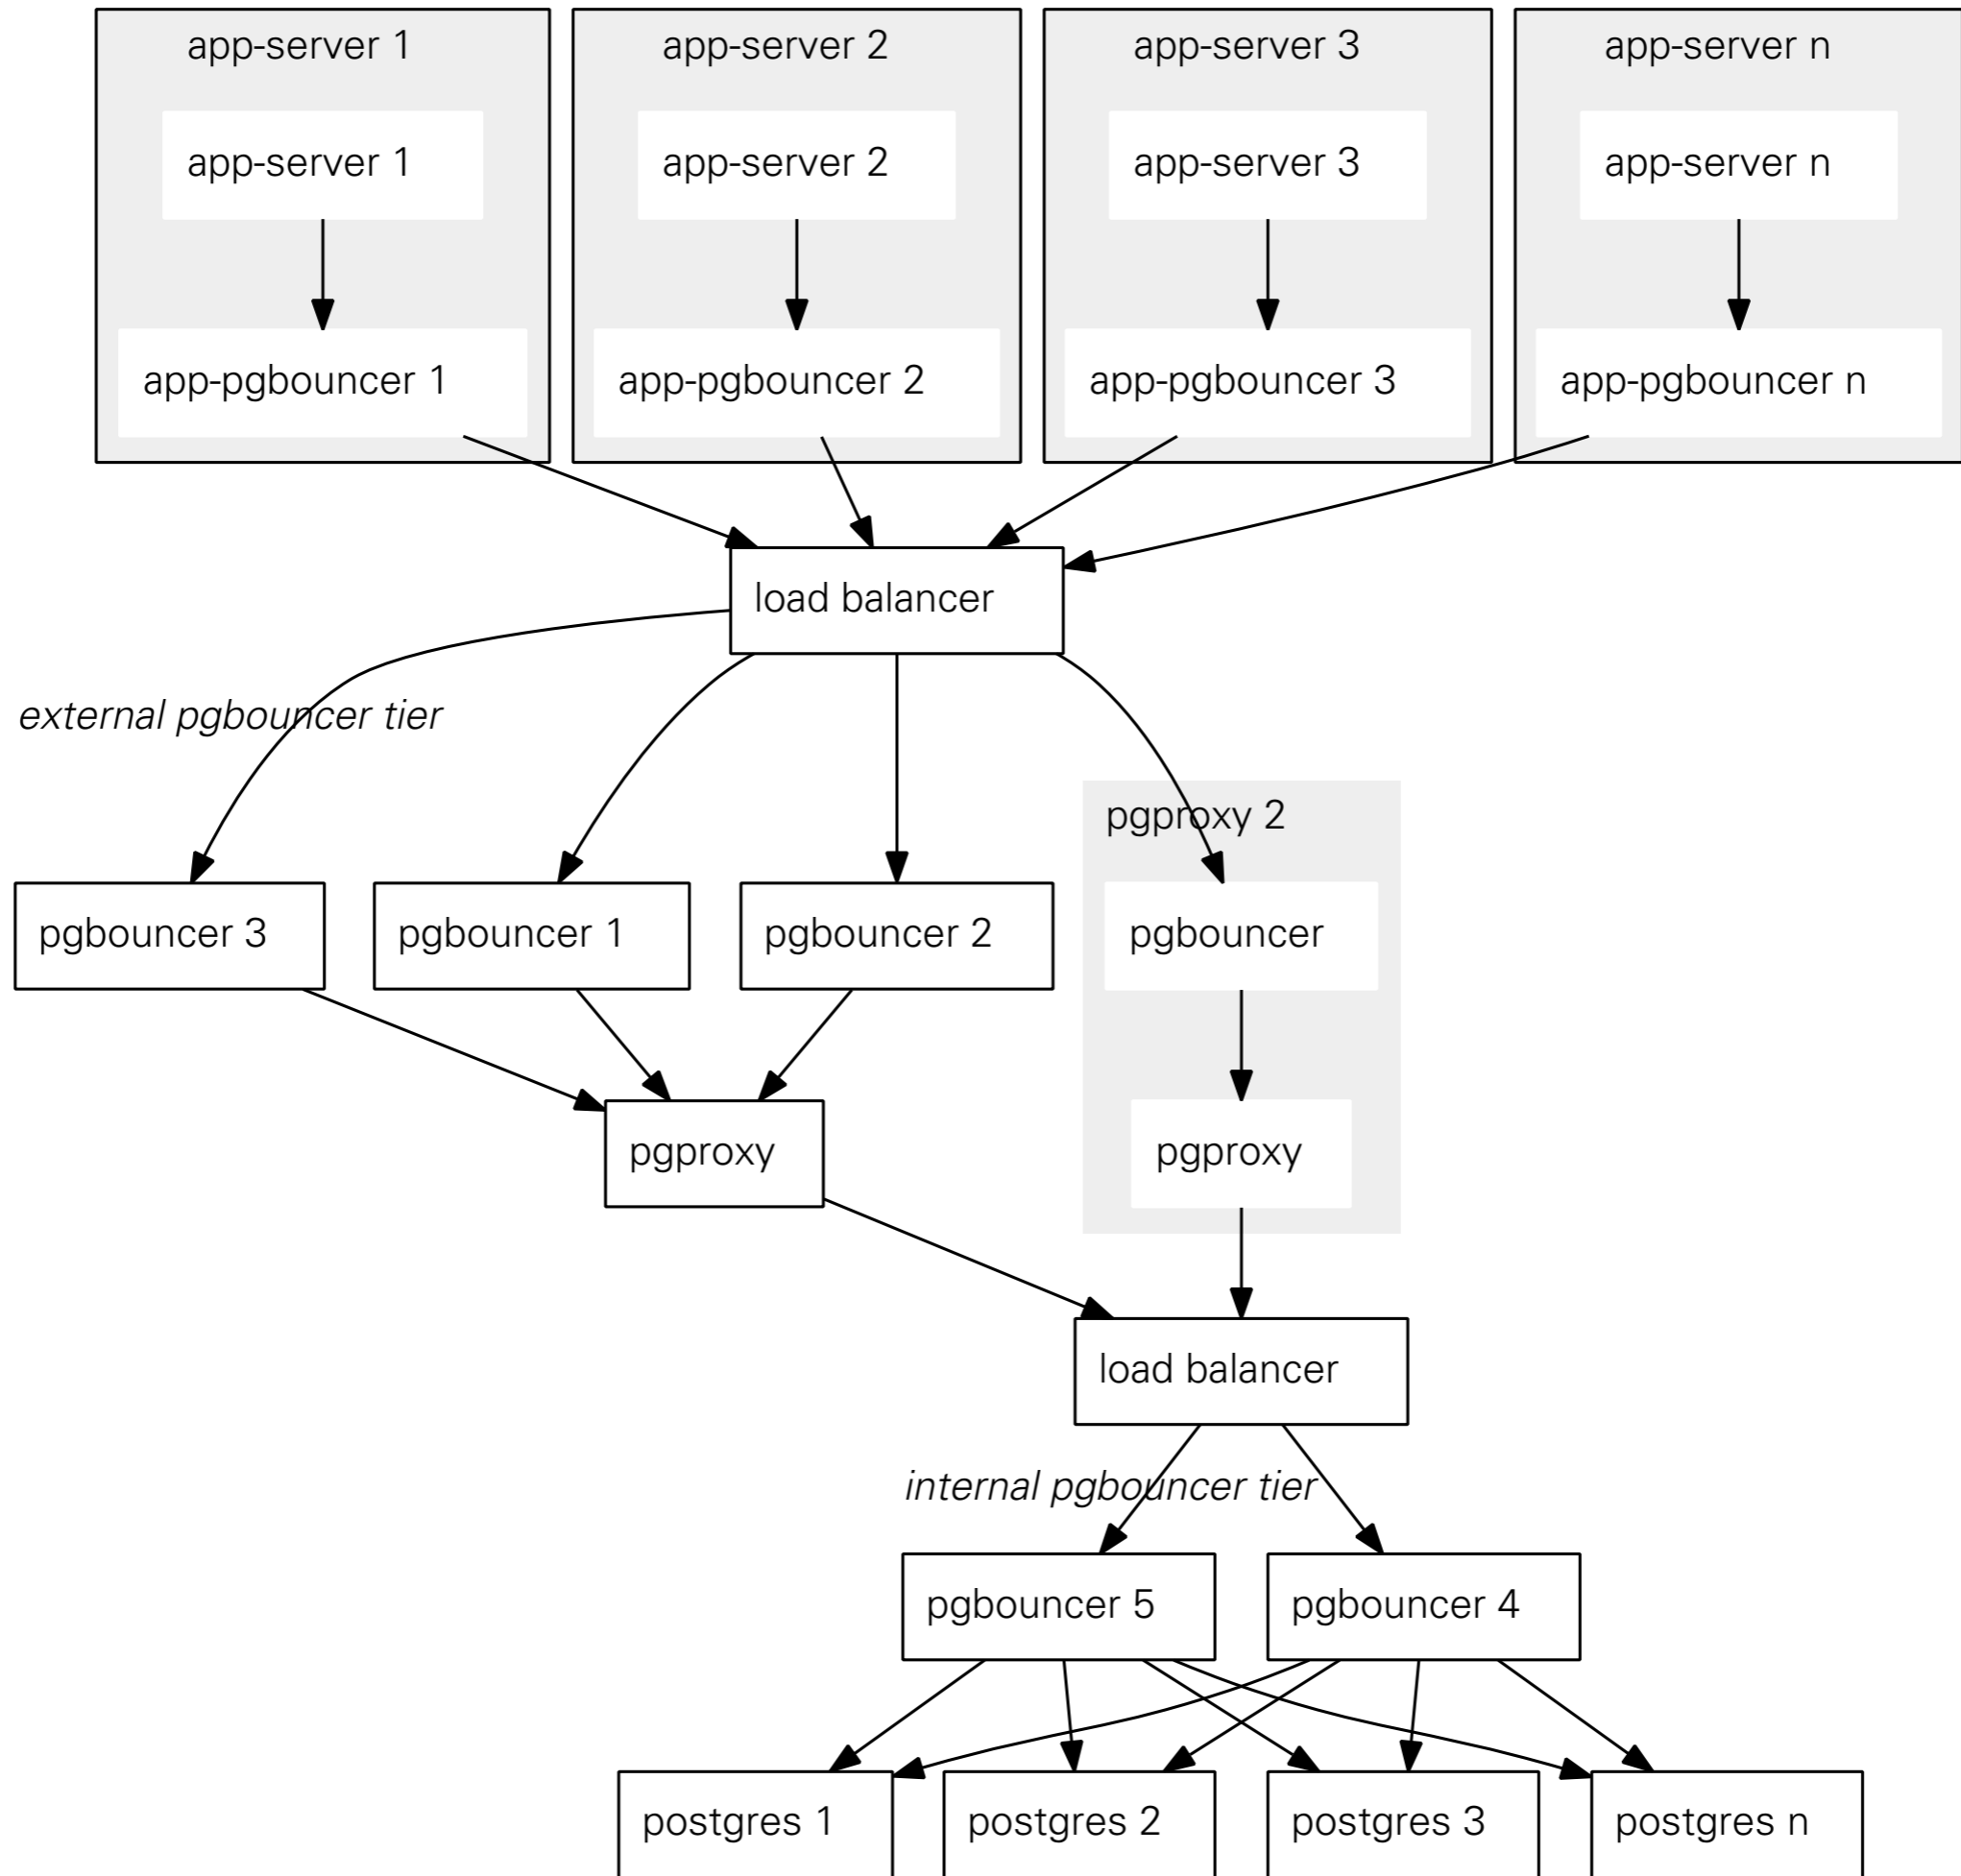


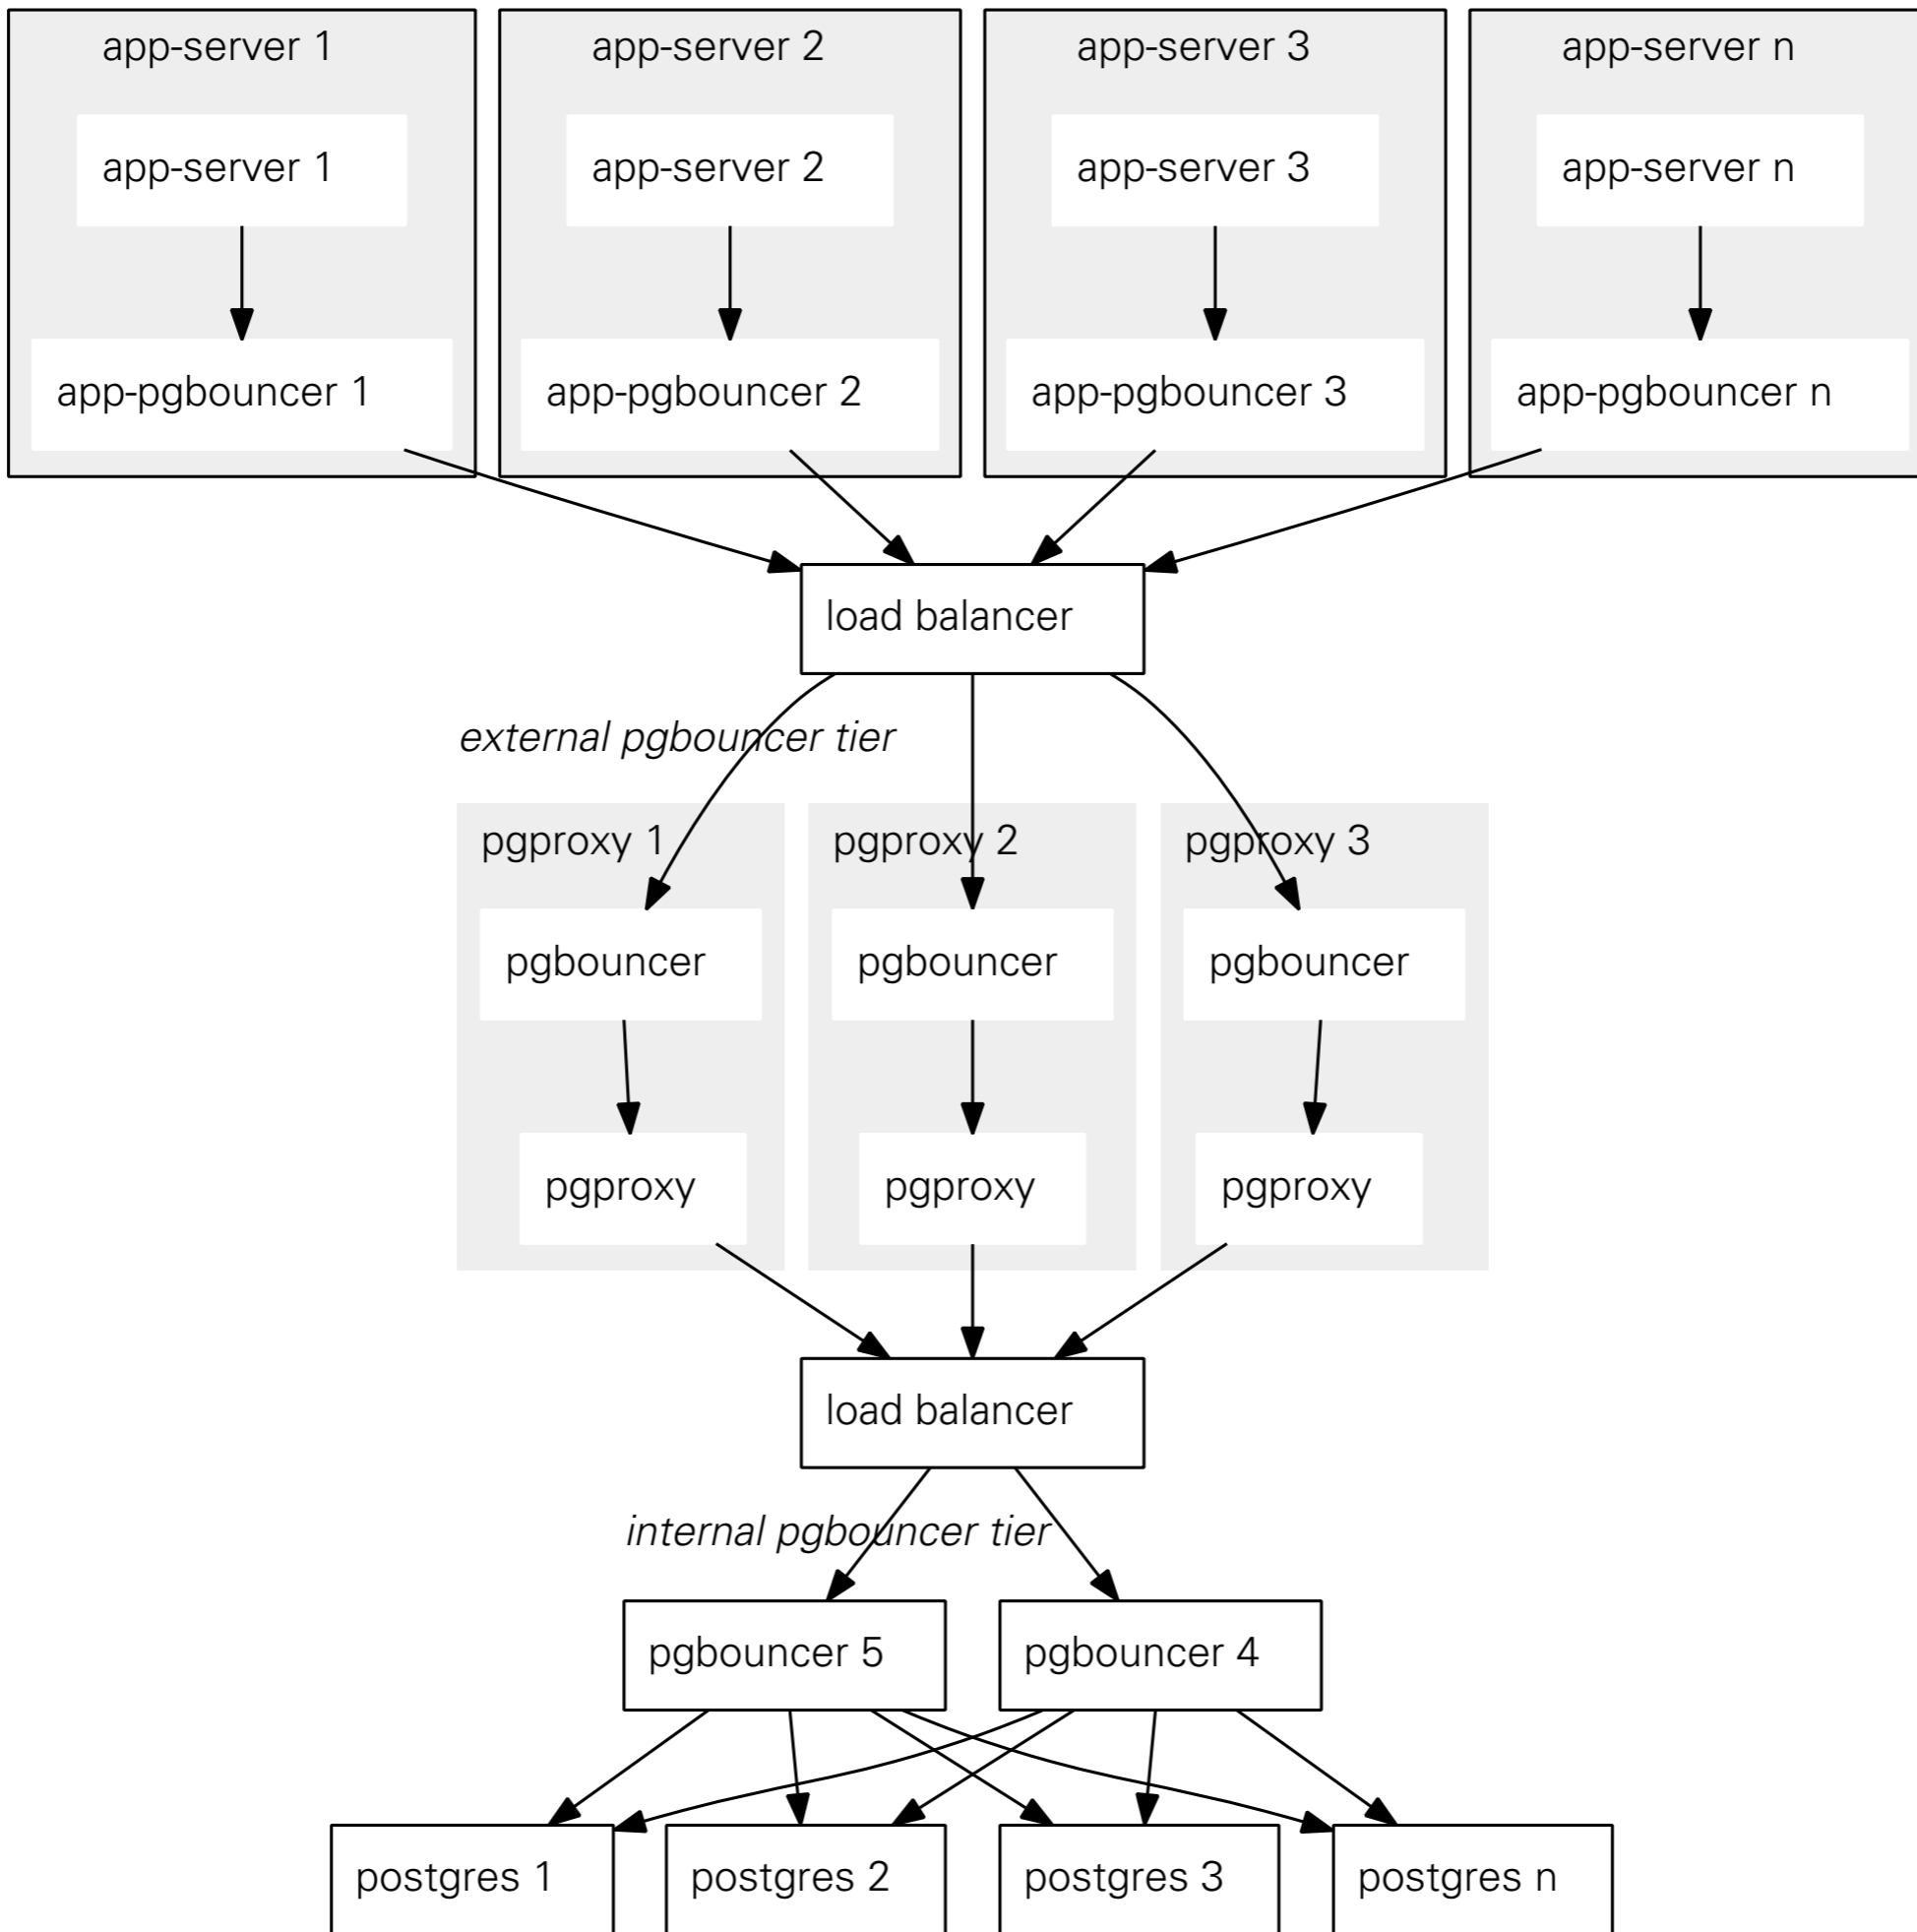
reduce TPS/server
less configuration











reduce TPS!

→ memcached

1 TB

get 140K/s, set 15K/s

pgfouine

connection pooling

→ pgbouncer

simplify interface

→ PL/Proxy

function API

partition

sharding

inter and intra server

roll off old data

asynchronous

→ message queues

28 servers avg 90% idle

464 cores

3.3 TB memory

3.8 TB on disk

35 TB total disk

15K avg TPS (> 27K)