

Lessons Learned Scaling a Social Network

michael.glaesemann@myyearbook.com

senior data architect

O'Reilly MySQL Conference & Expo 2011
2011-04-13

Who?

日本

UTF-8

PostgreSQL

community

Codd

Why?

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

2007

100M HTTP requests/month

1 PostgreSQL server

1 Apache/PHP server

static content servers

1 outage per night

2008

2.5G page views

30 PostgreSQL servers

120 Apache/PHP servers

8 ActiveMQ message brokers

99.84% uptime

2010

2.5G HTTP requests/month

Database servers

connection poolers, routers

Web servers

Apache/PHP, Tornado/Python, lighttpd

Message brokers

ActiveMQ, RabbitMQ

Monitoring, R&D, other servers

rapid
growth

resources

1

API

bottleneck

time

Relative data access latencies

1	CPU Register
1–2	L1 cache
6–10	L2 cache
25–50	main memory (RAM)
10,000,000	hard disk
100,000,000	LAN
1,000,000,000	WAN

fsync

memory

more activity?

views/month 2007 100M → 2009 1.5G

more TPS

more servers

more connections

more configuration

more pain!

reduce TPS!

→ memcached

1.2 TB

get 140K/s, set 15K/s

pgfouine


Skype

PL/Proxy


pgBouncer

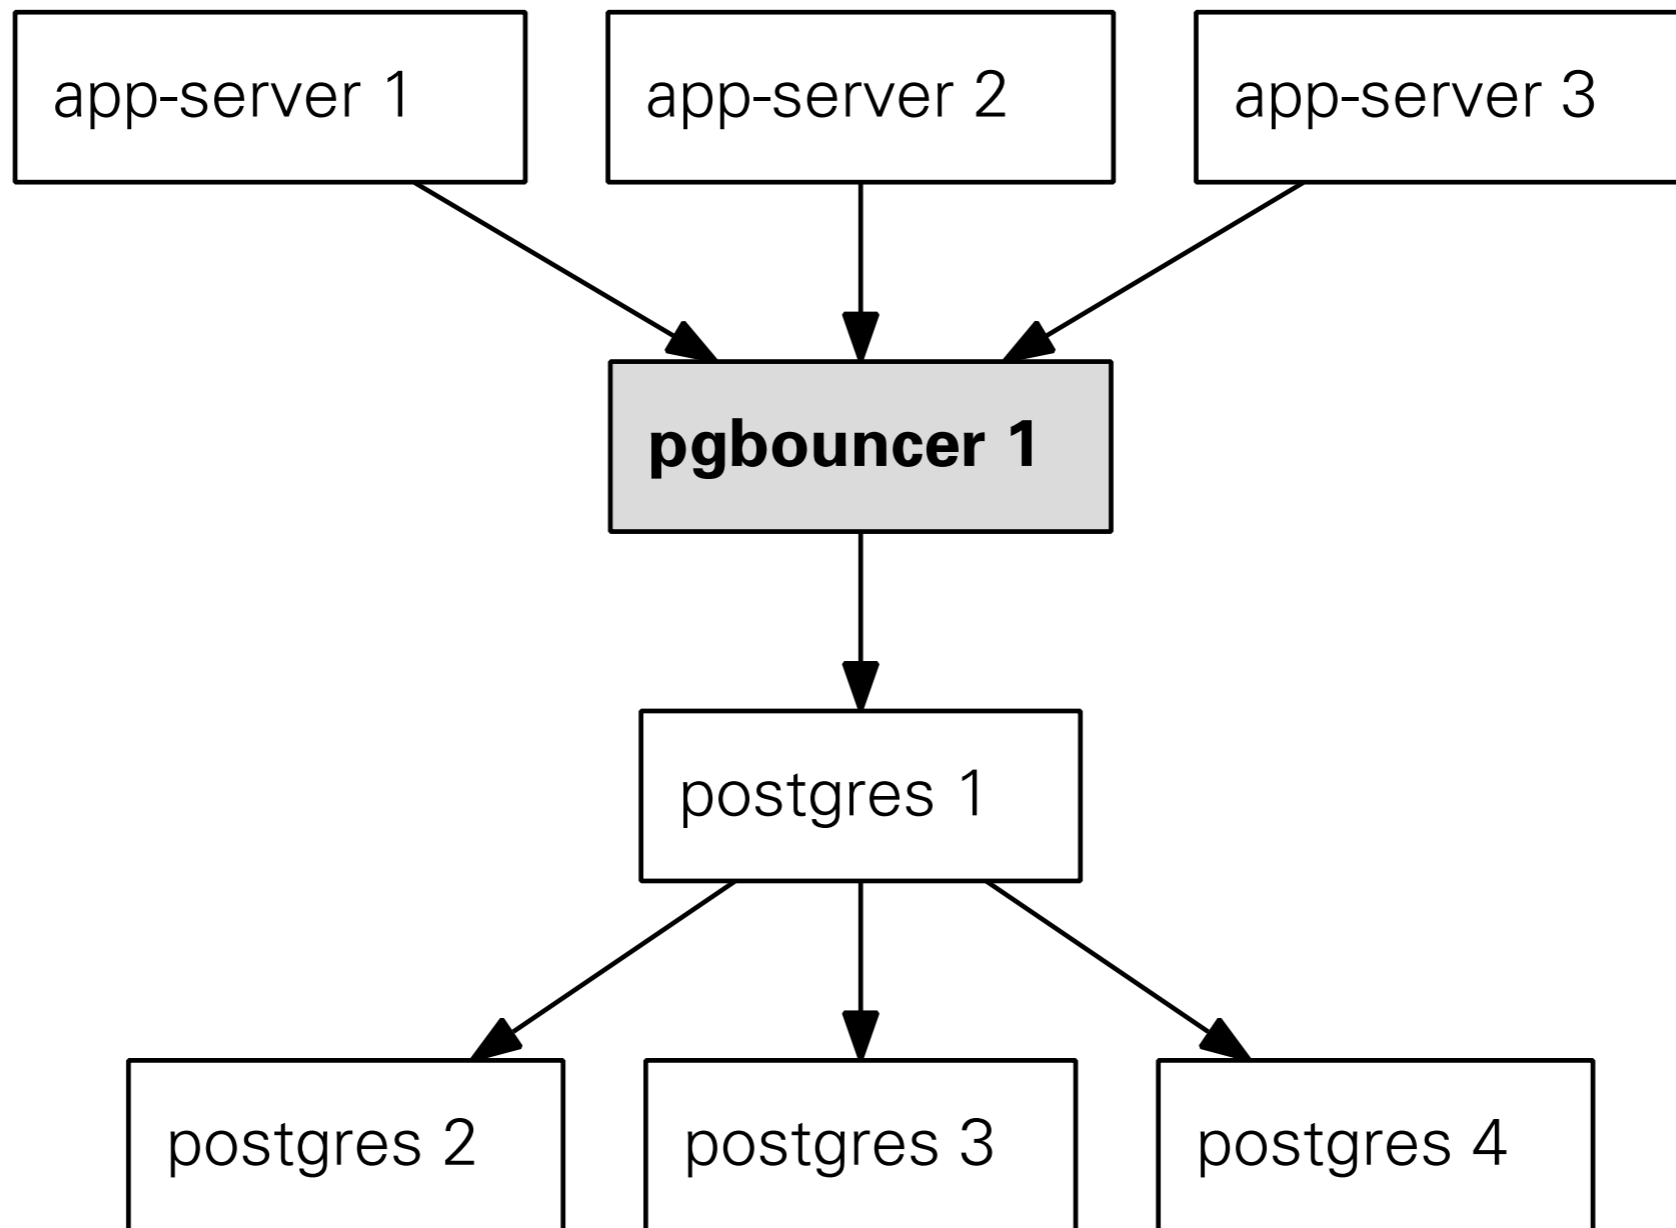
less configuration

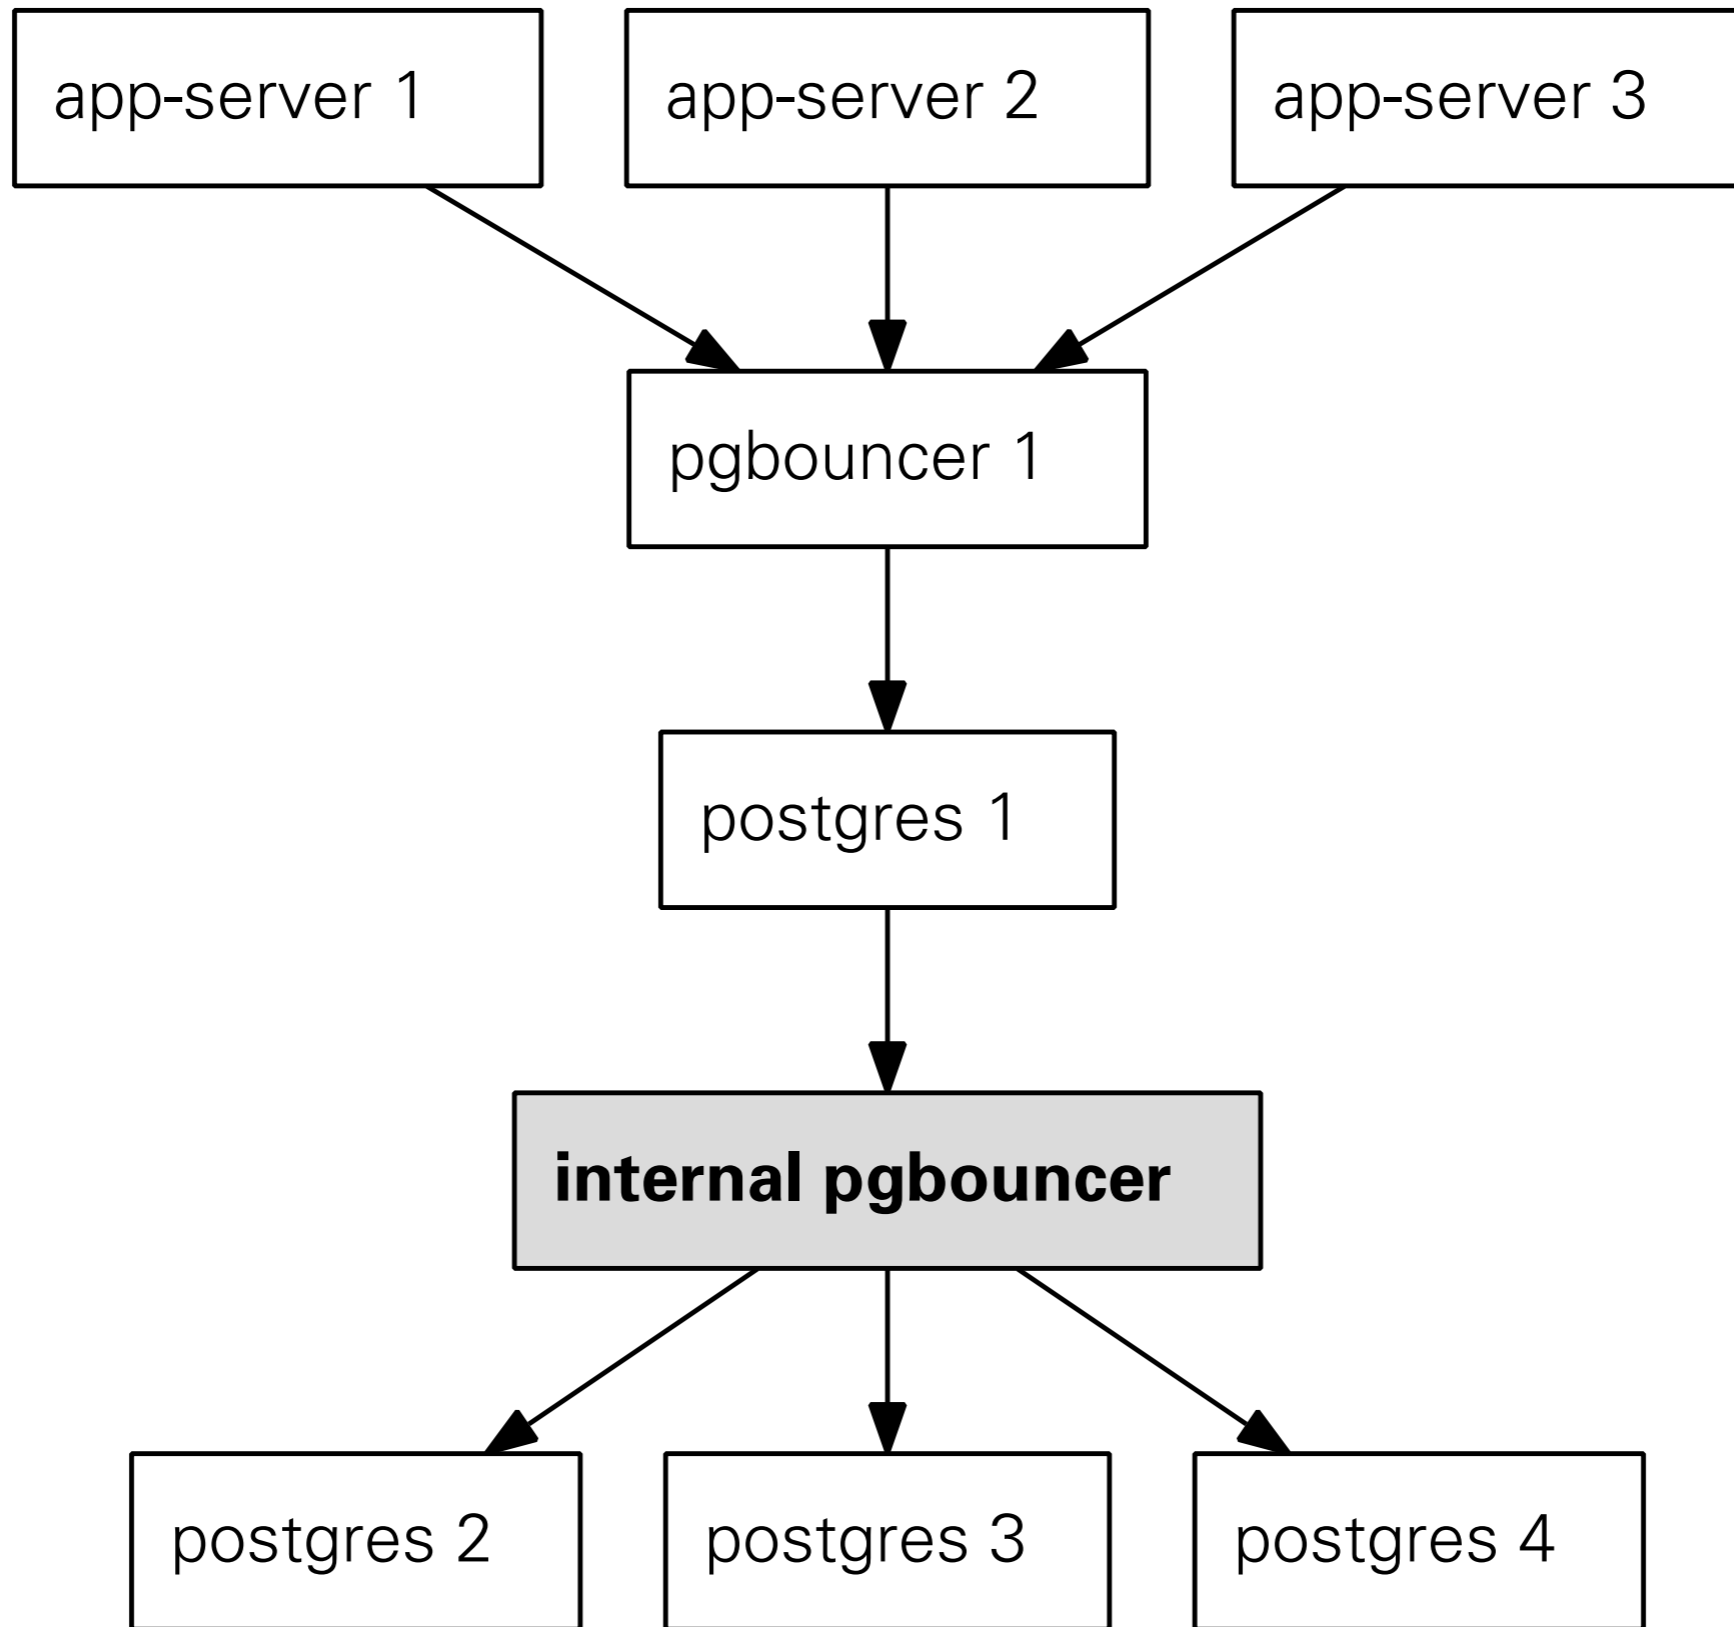
fewer connections




fewer connections!




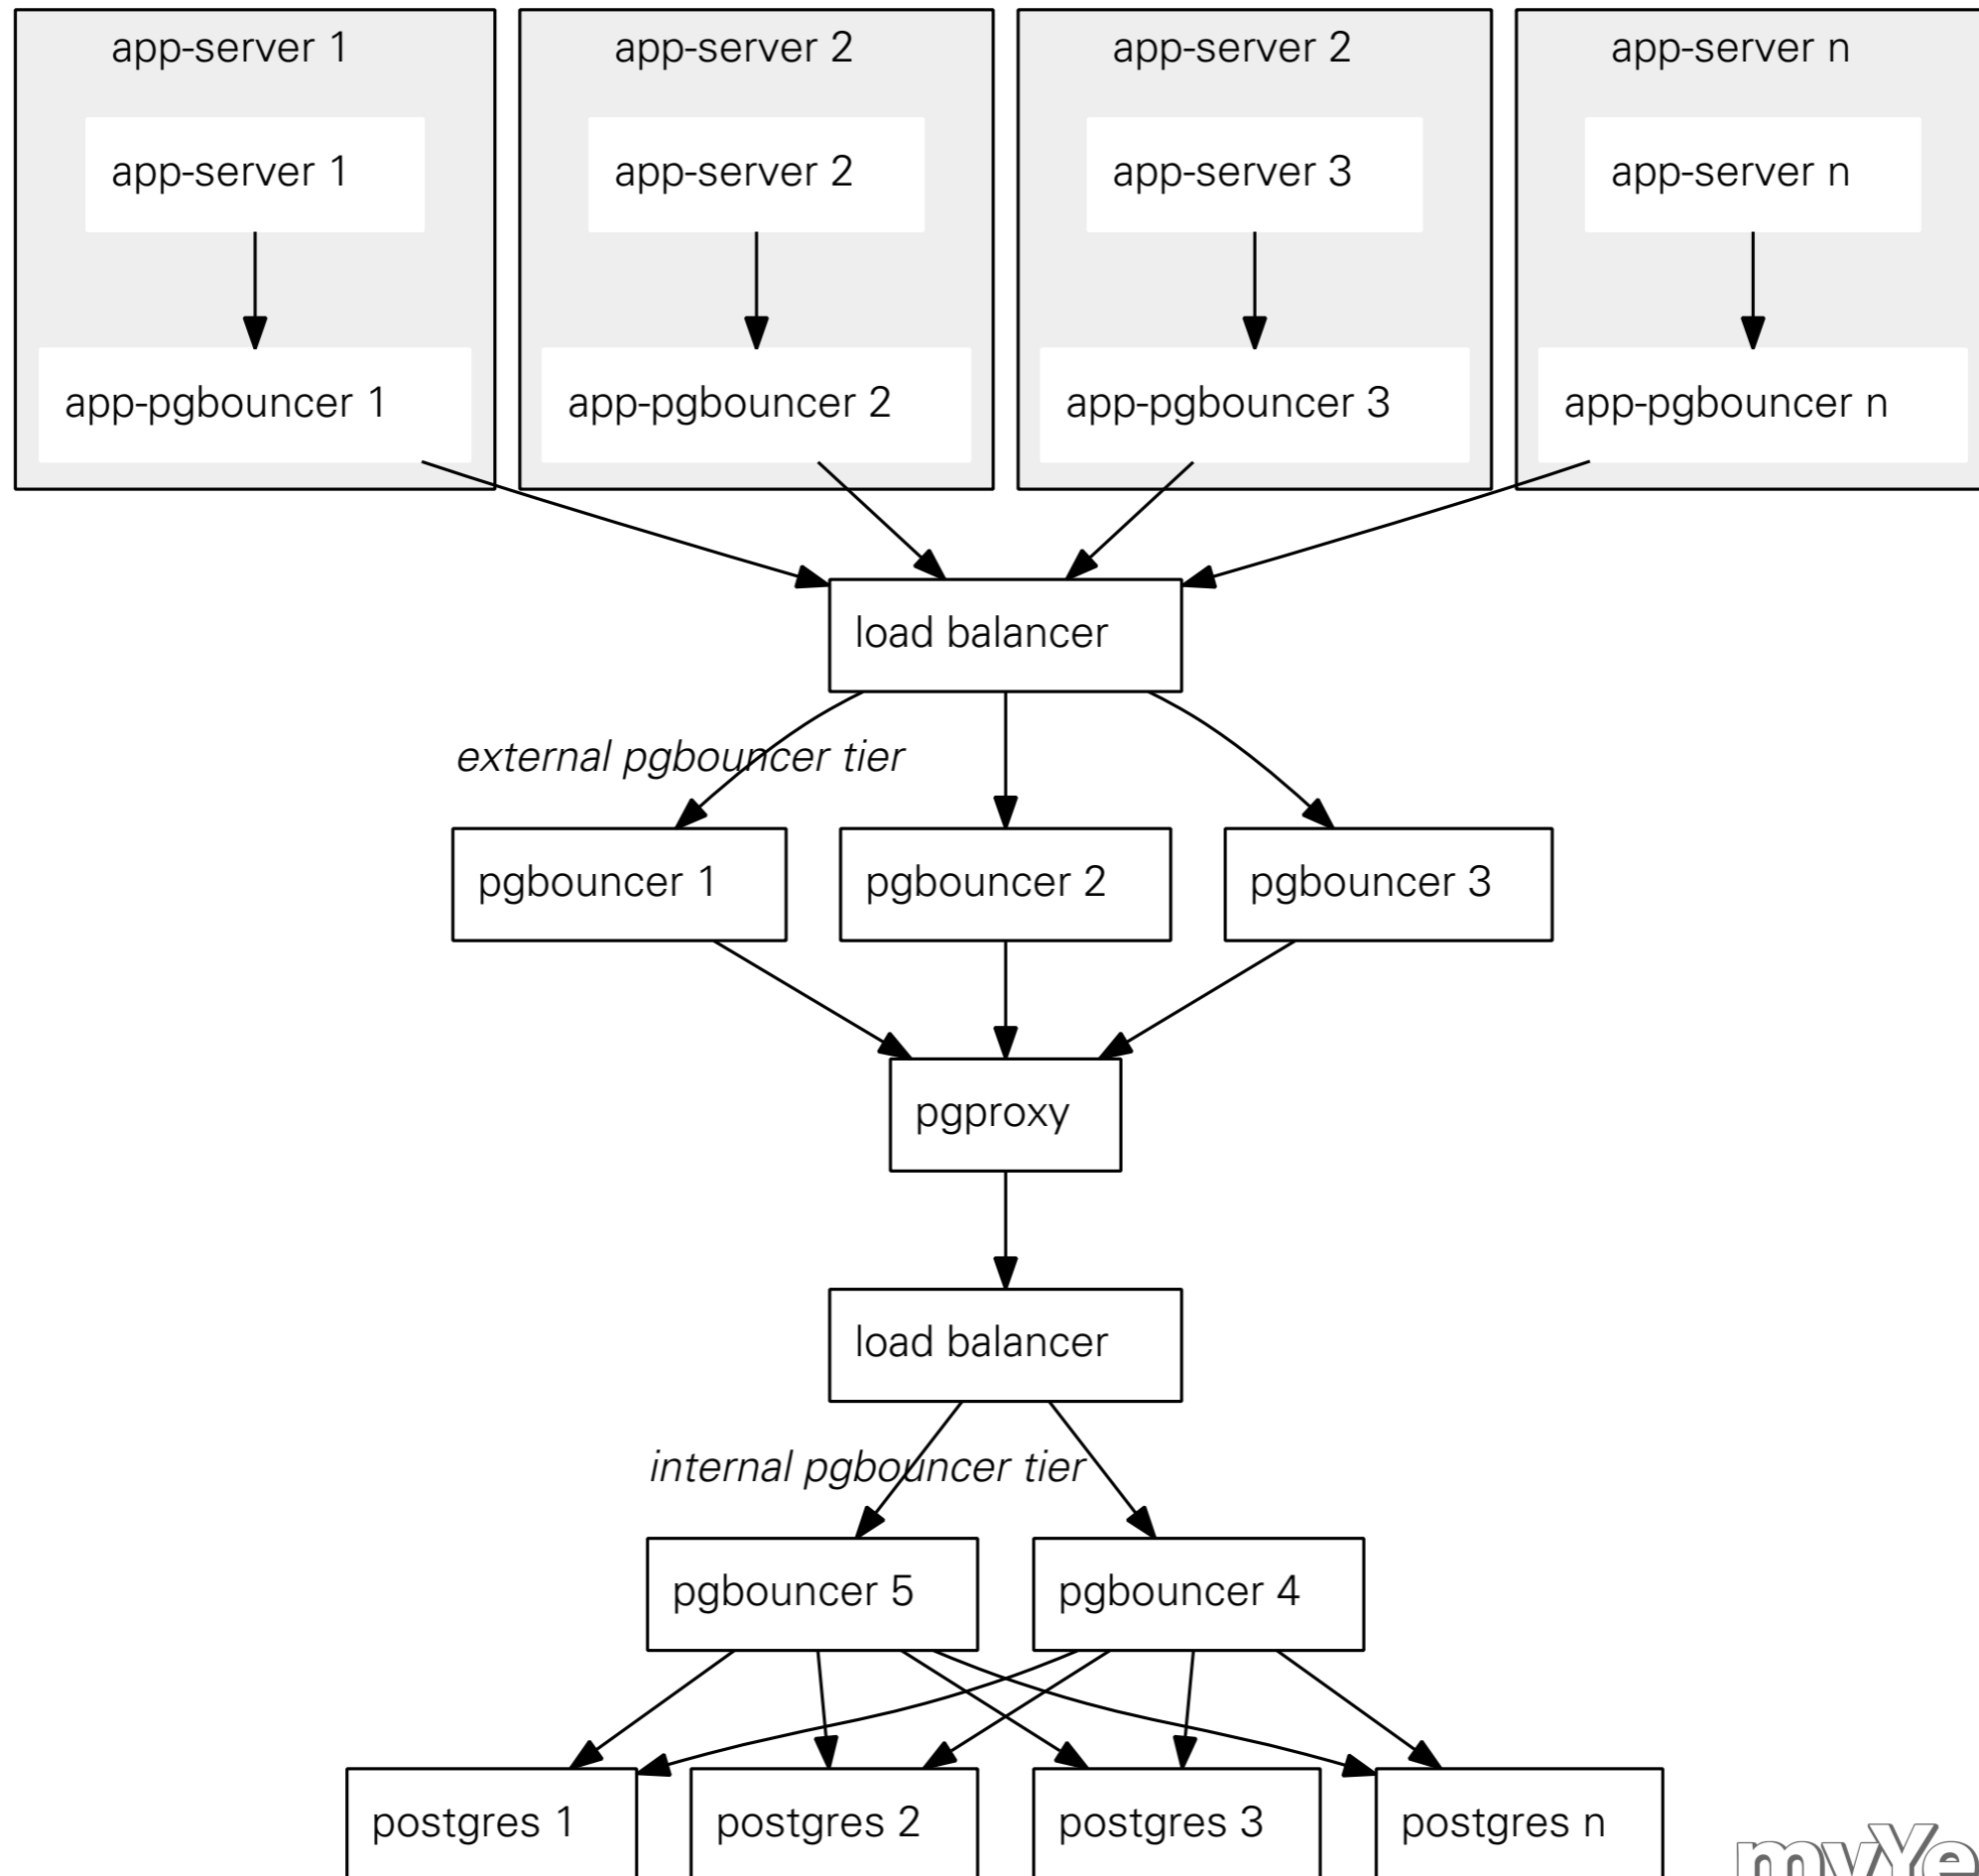


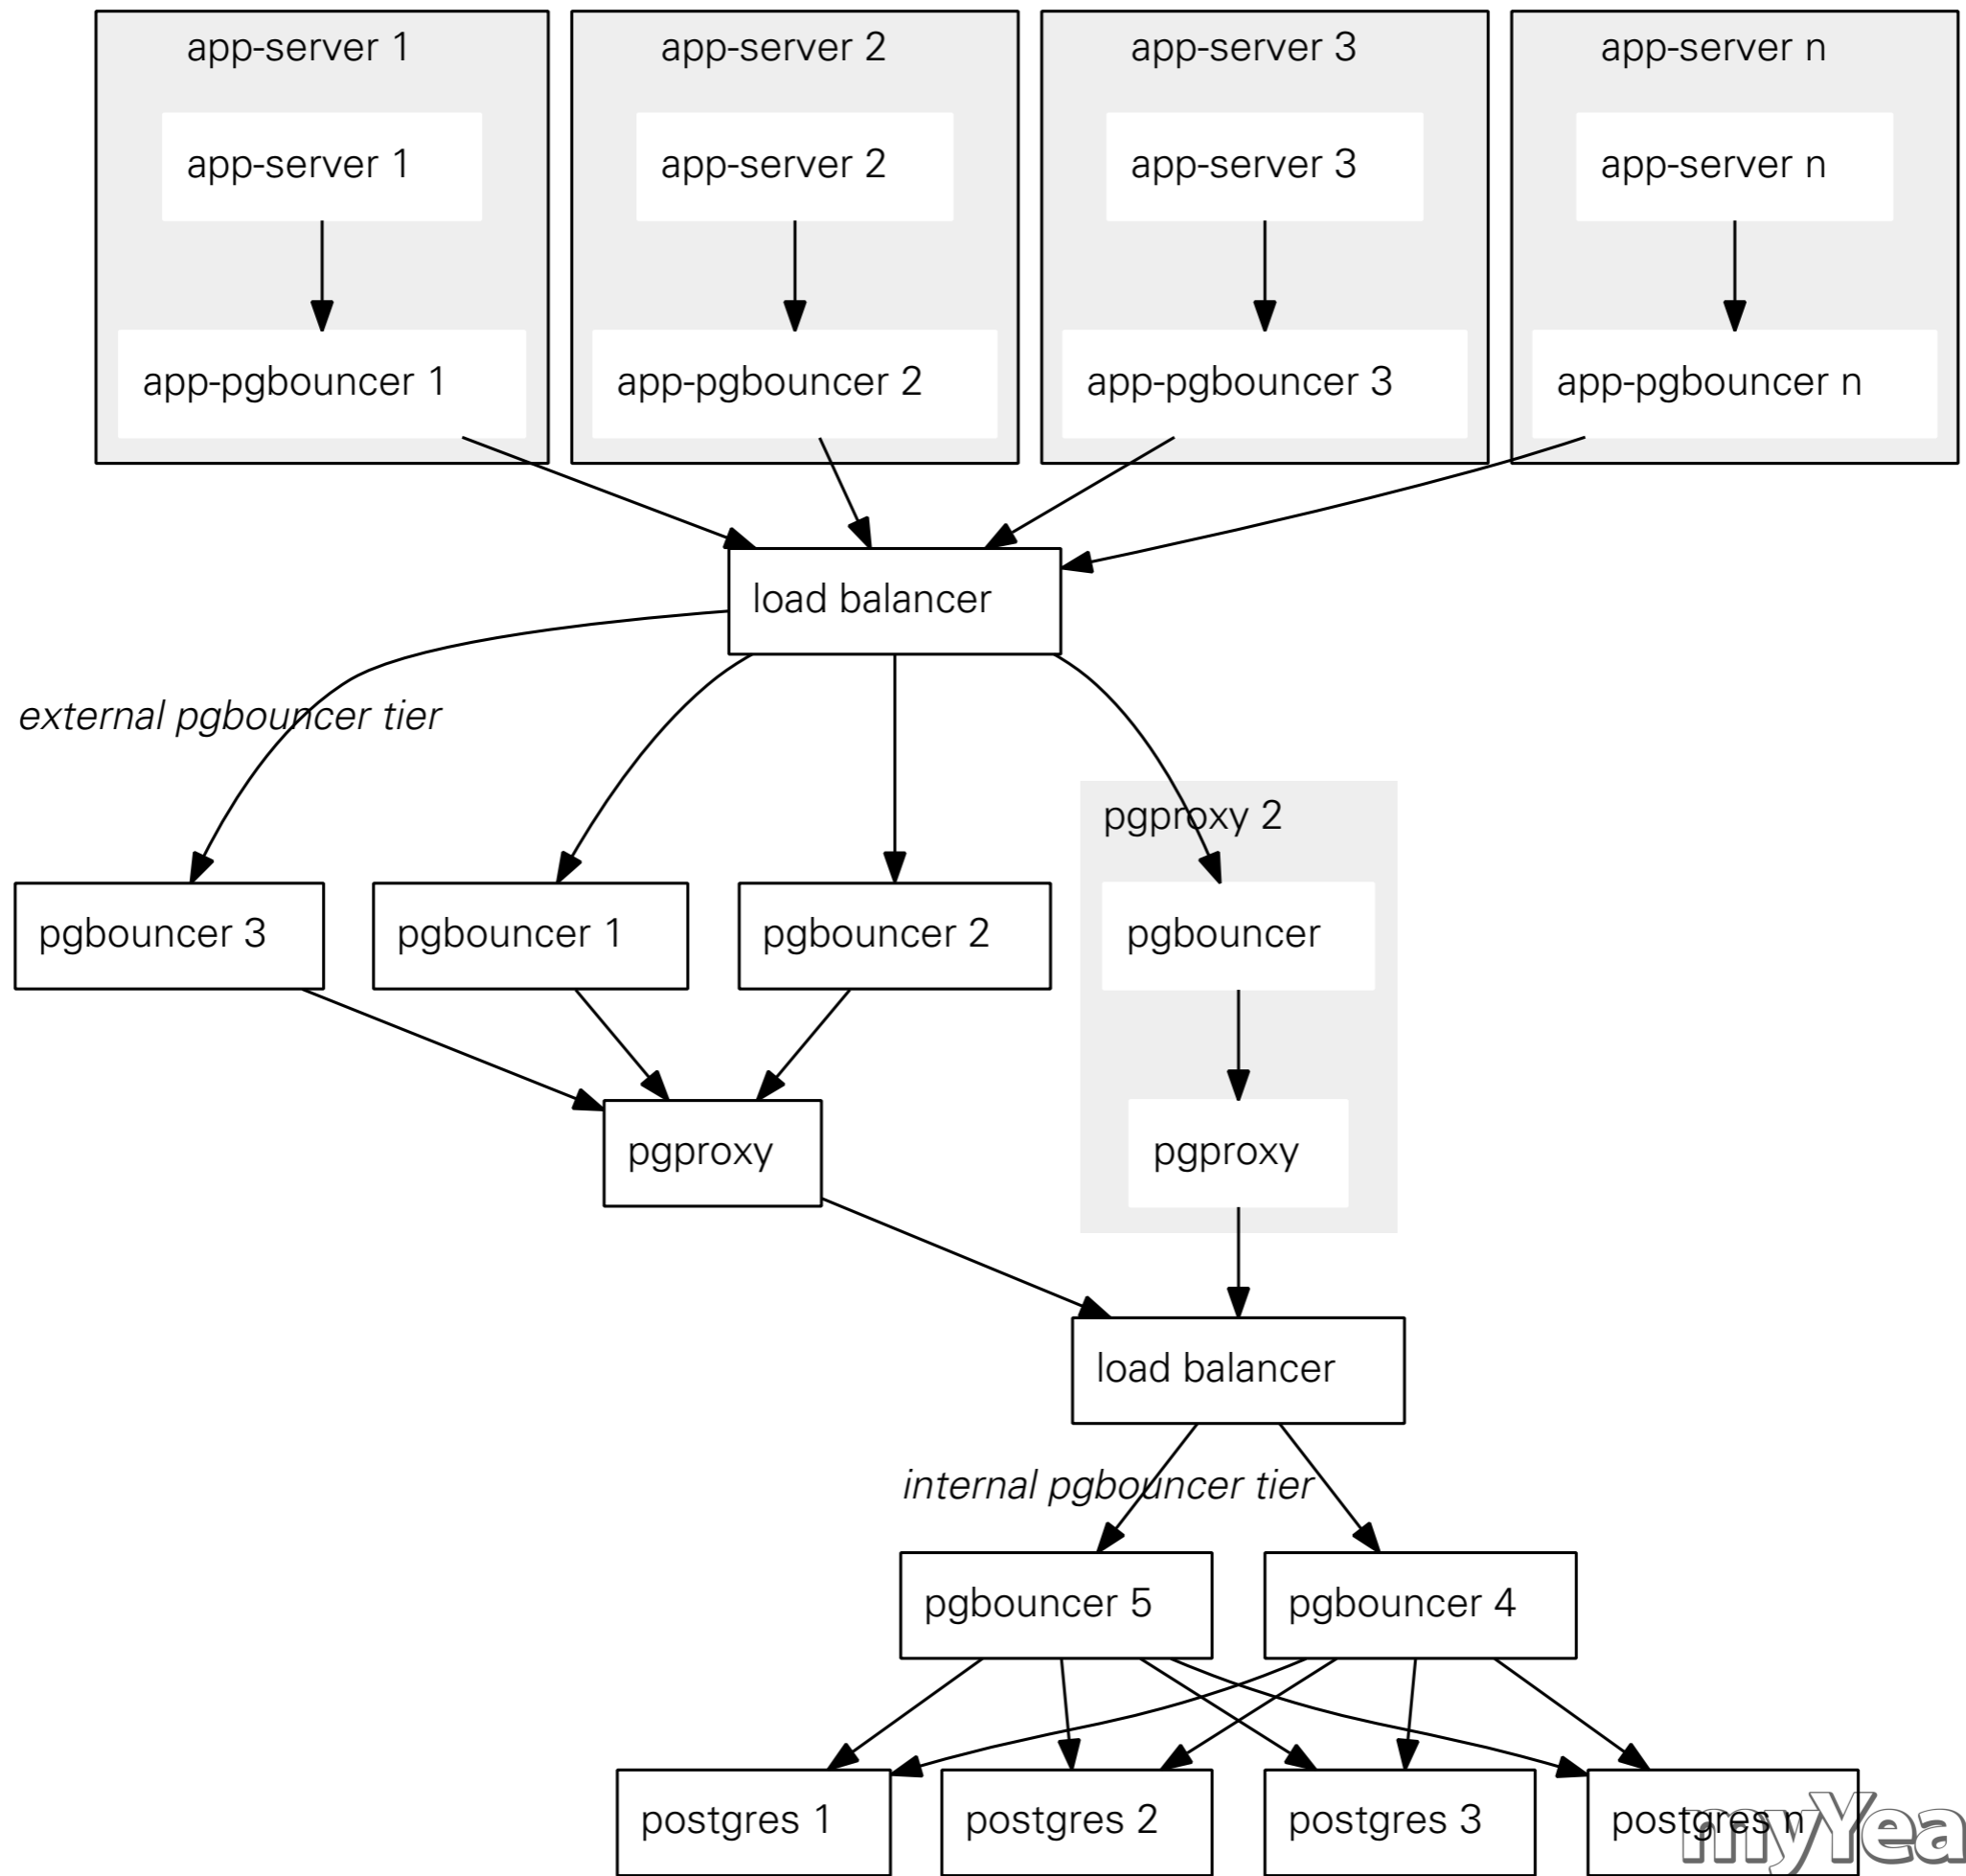



reduce TPS/server
less configuration











1

SOA

REST

DB API

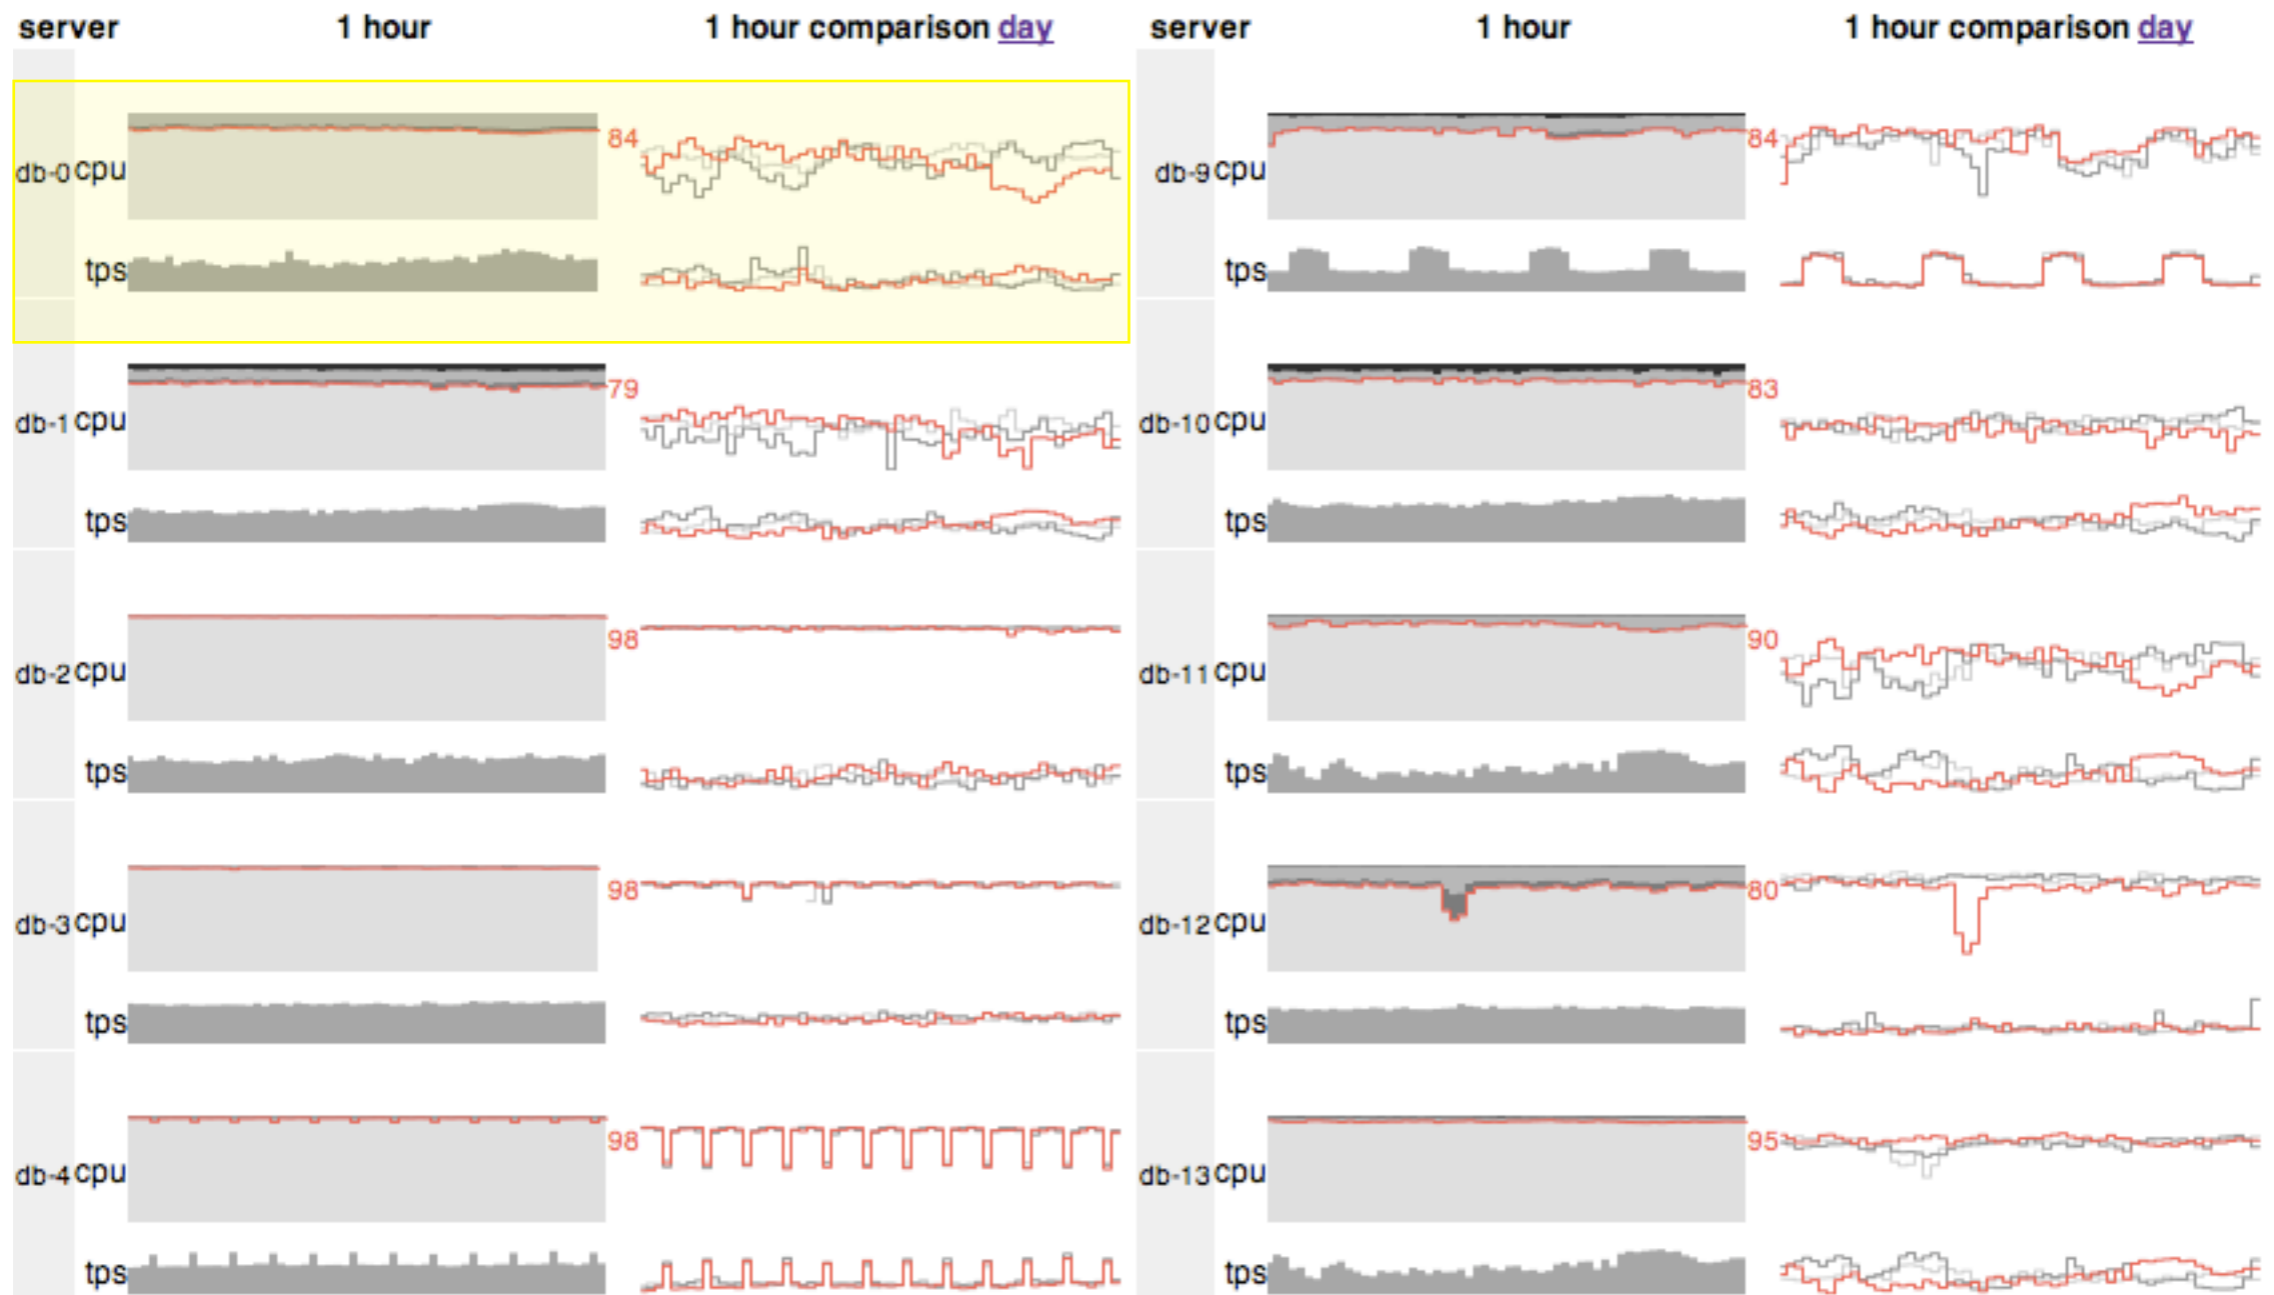
death

failure is
a choice

network

measure

overview



learn

logical
versus
physical


thought

word

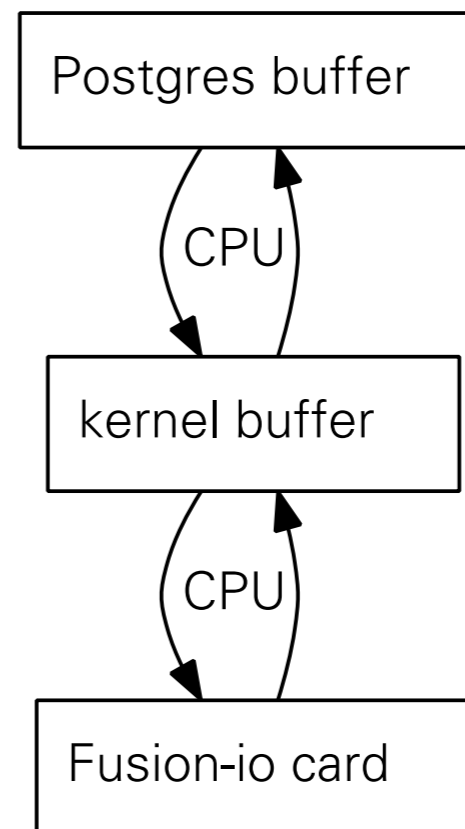
deed

Fusion-io

direct IO



buffered IO



improve

automate

test

Lessons Learned Scaling a Social Network

michael.glaesemann@myyearbook.com

Senior Data Architect

O'Reilly MySQL Conference & Expo 2011
2011-04-13

partition

sharding

inter and intra server

roll off old data