

# PostgreSQL 在 Qunar 的演进

wd@qunar

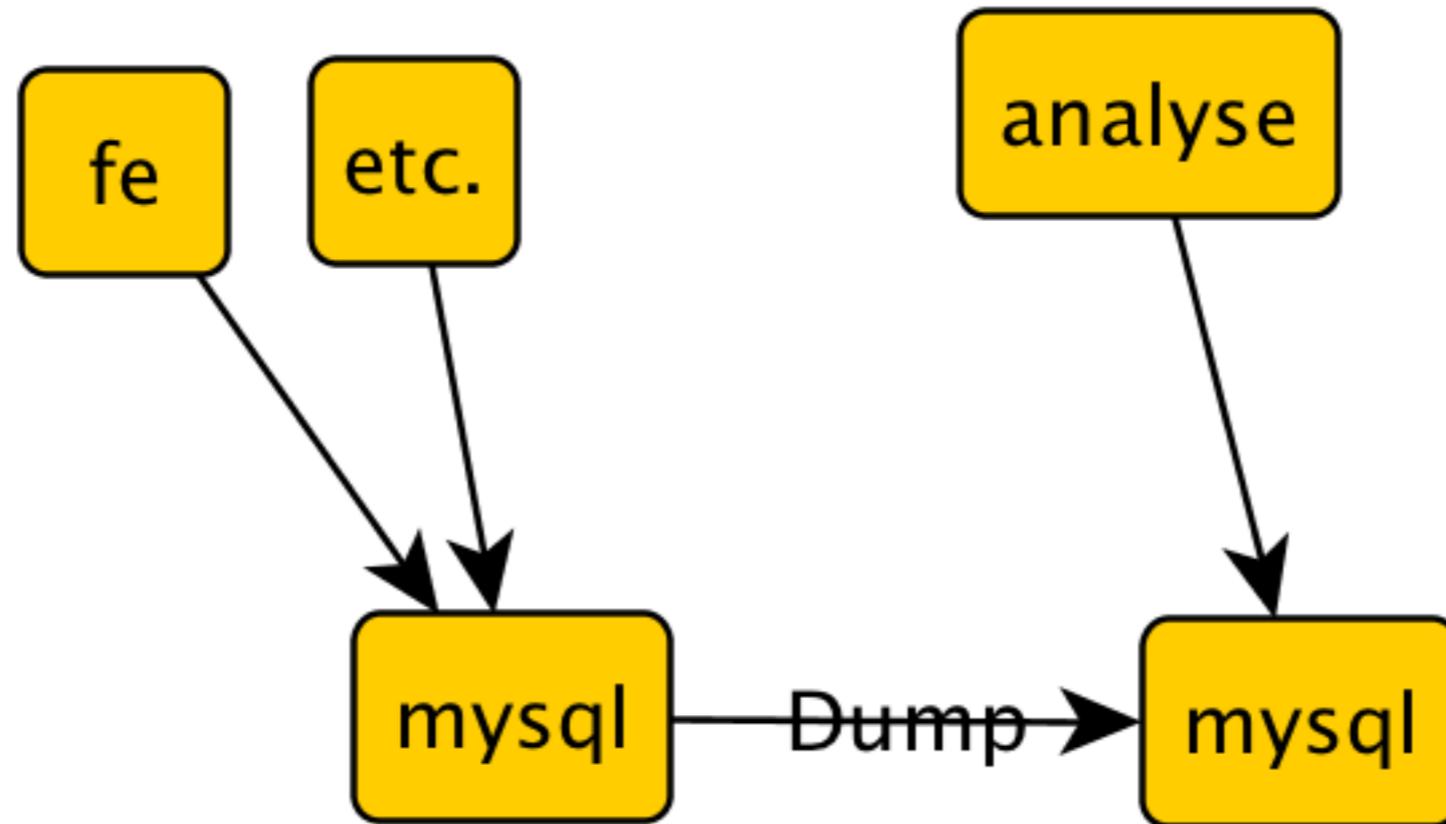
# TOC

- About me
- PostgreSQL 在 Qunar 的使用
- 针对 PostgreSQL 做的一些运维工作
- 接下来要做的事情

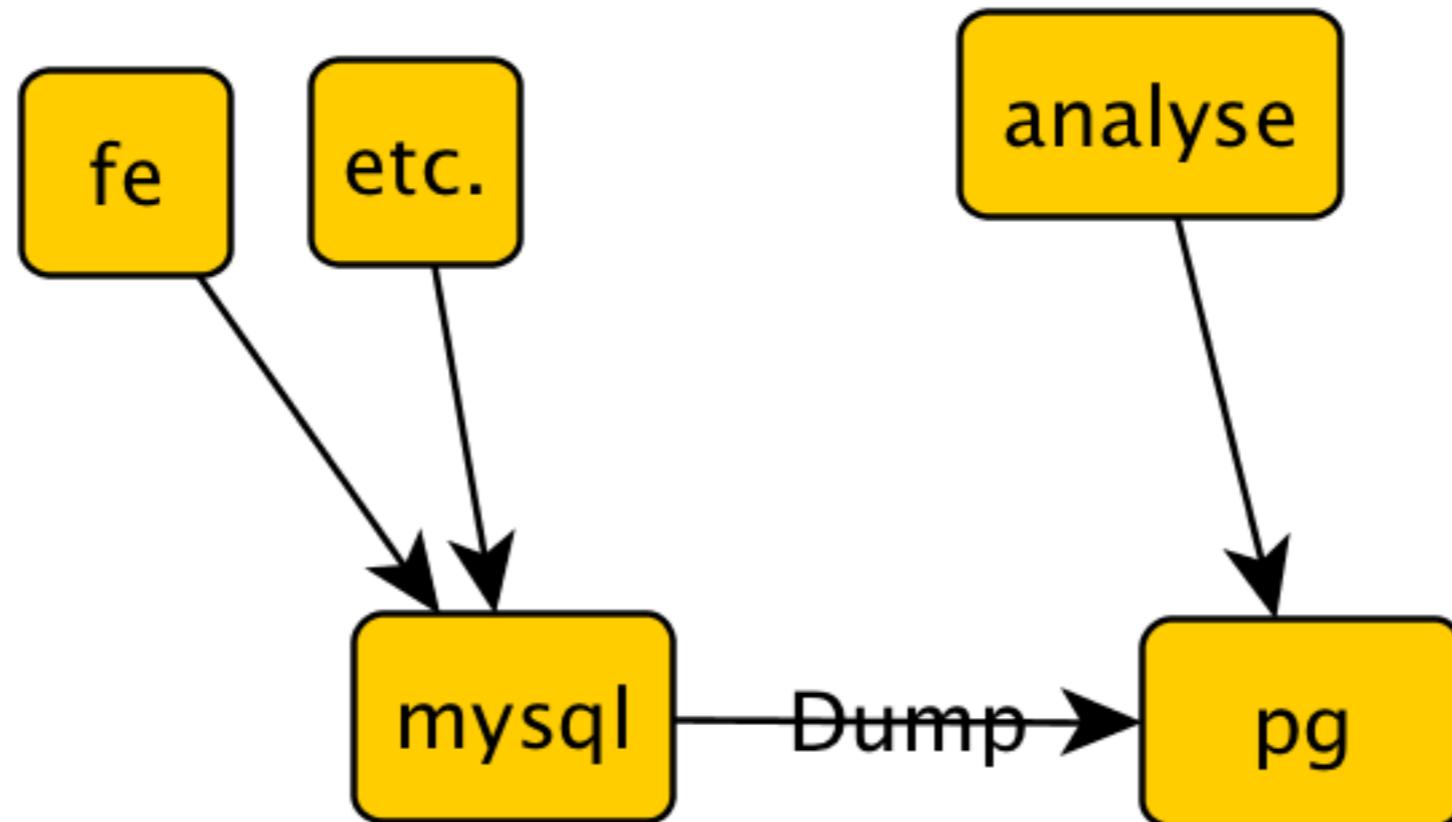
# About me

- 王冬(wd)
- PostgreSQL DBA
- 旅游度假部门移动端

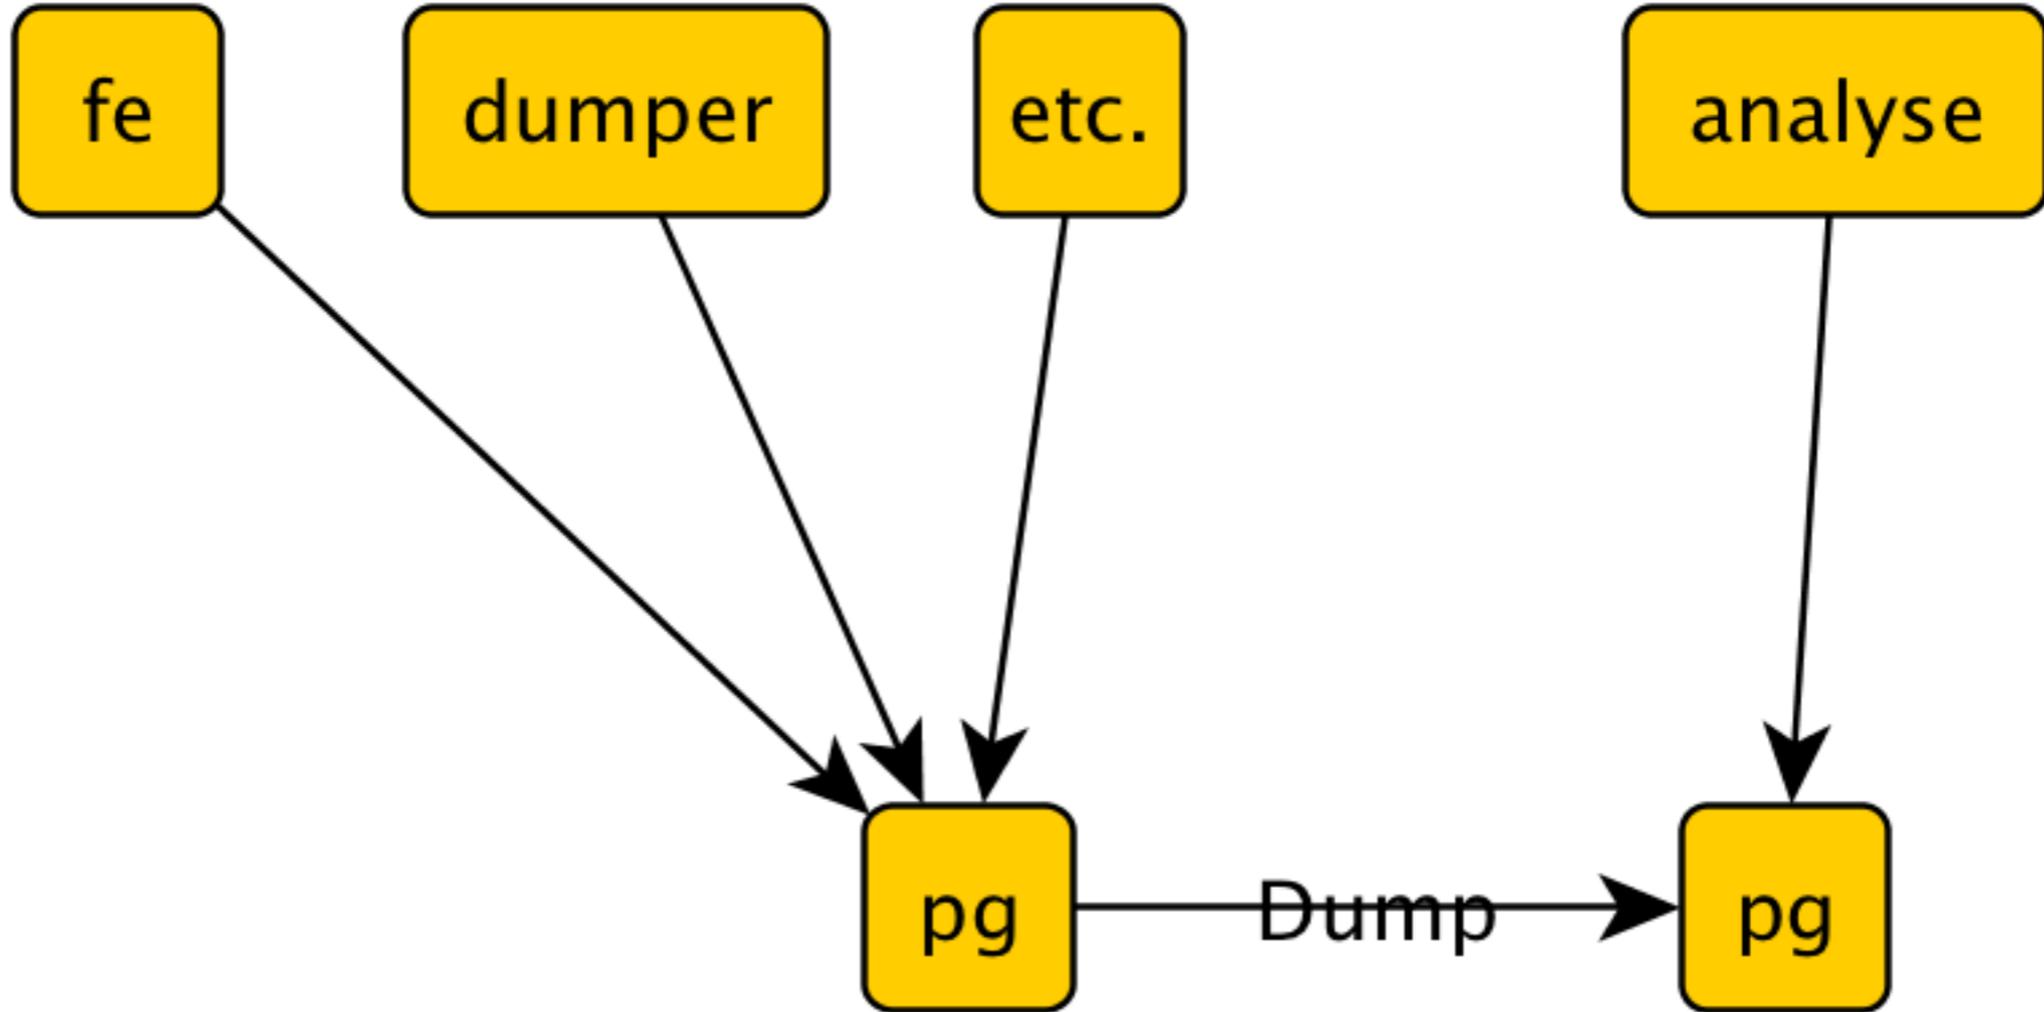
# 第一版结构



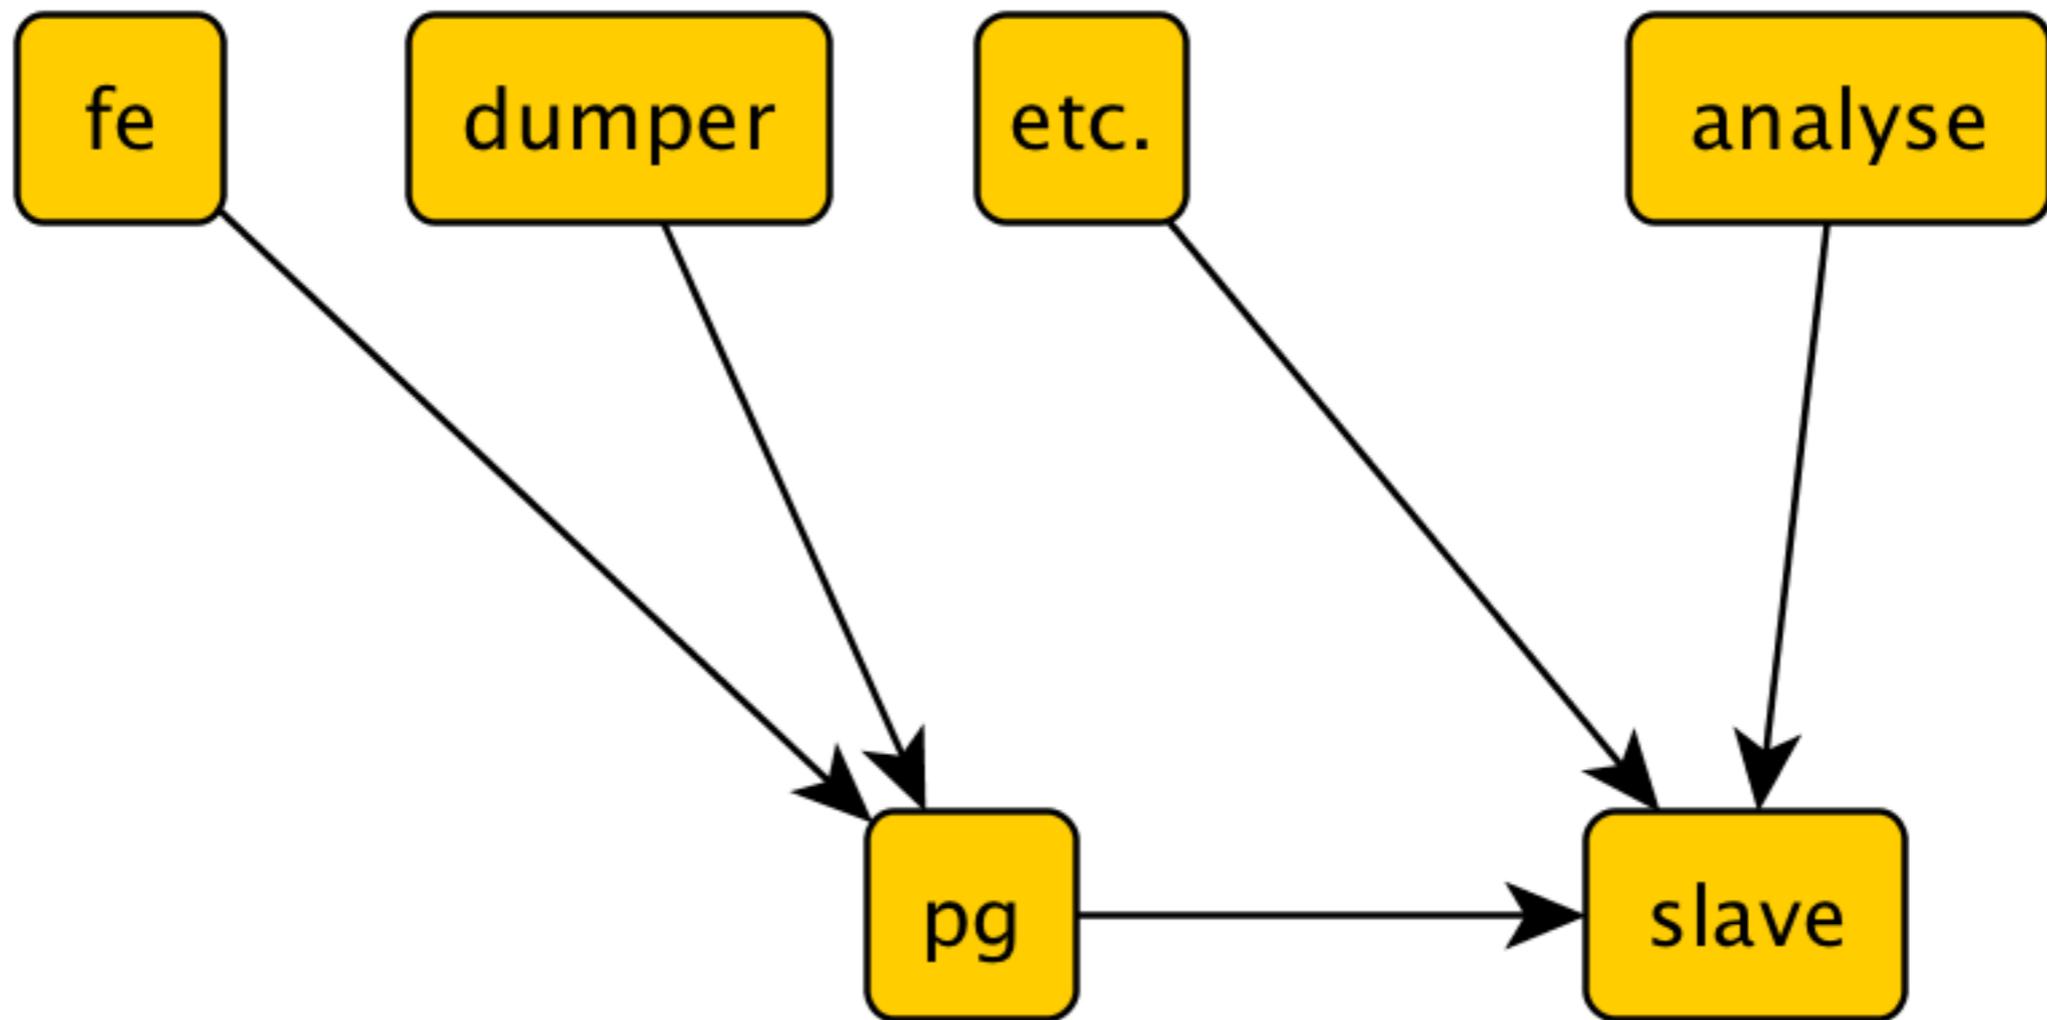
# 改进



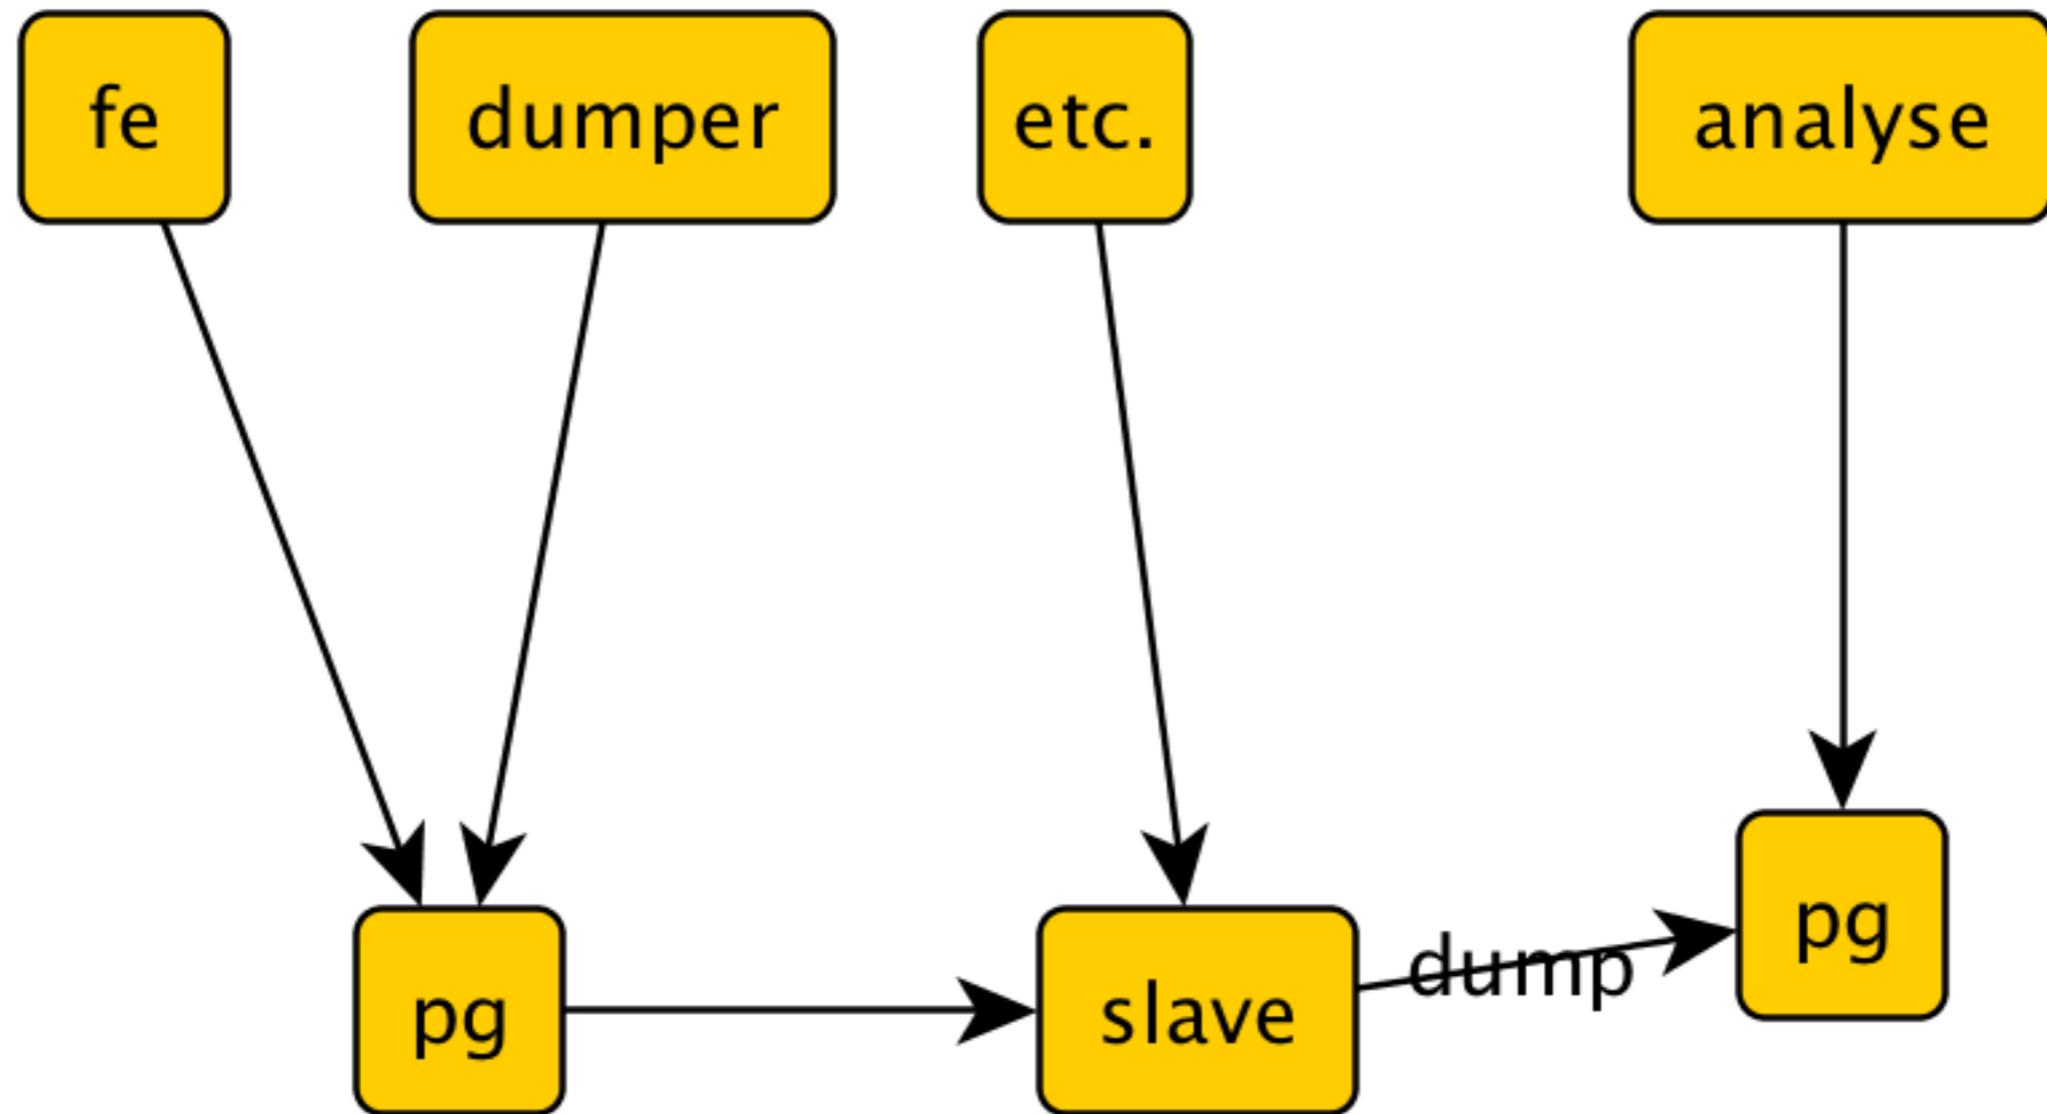
# 改进



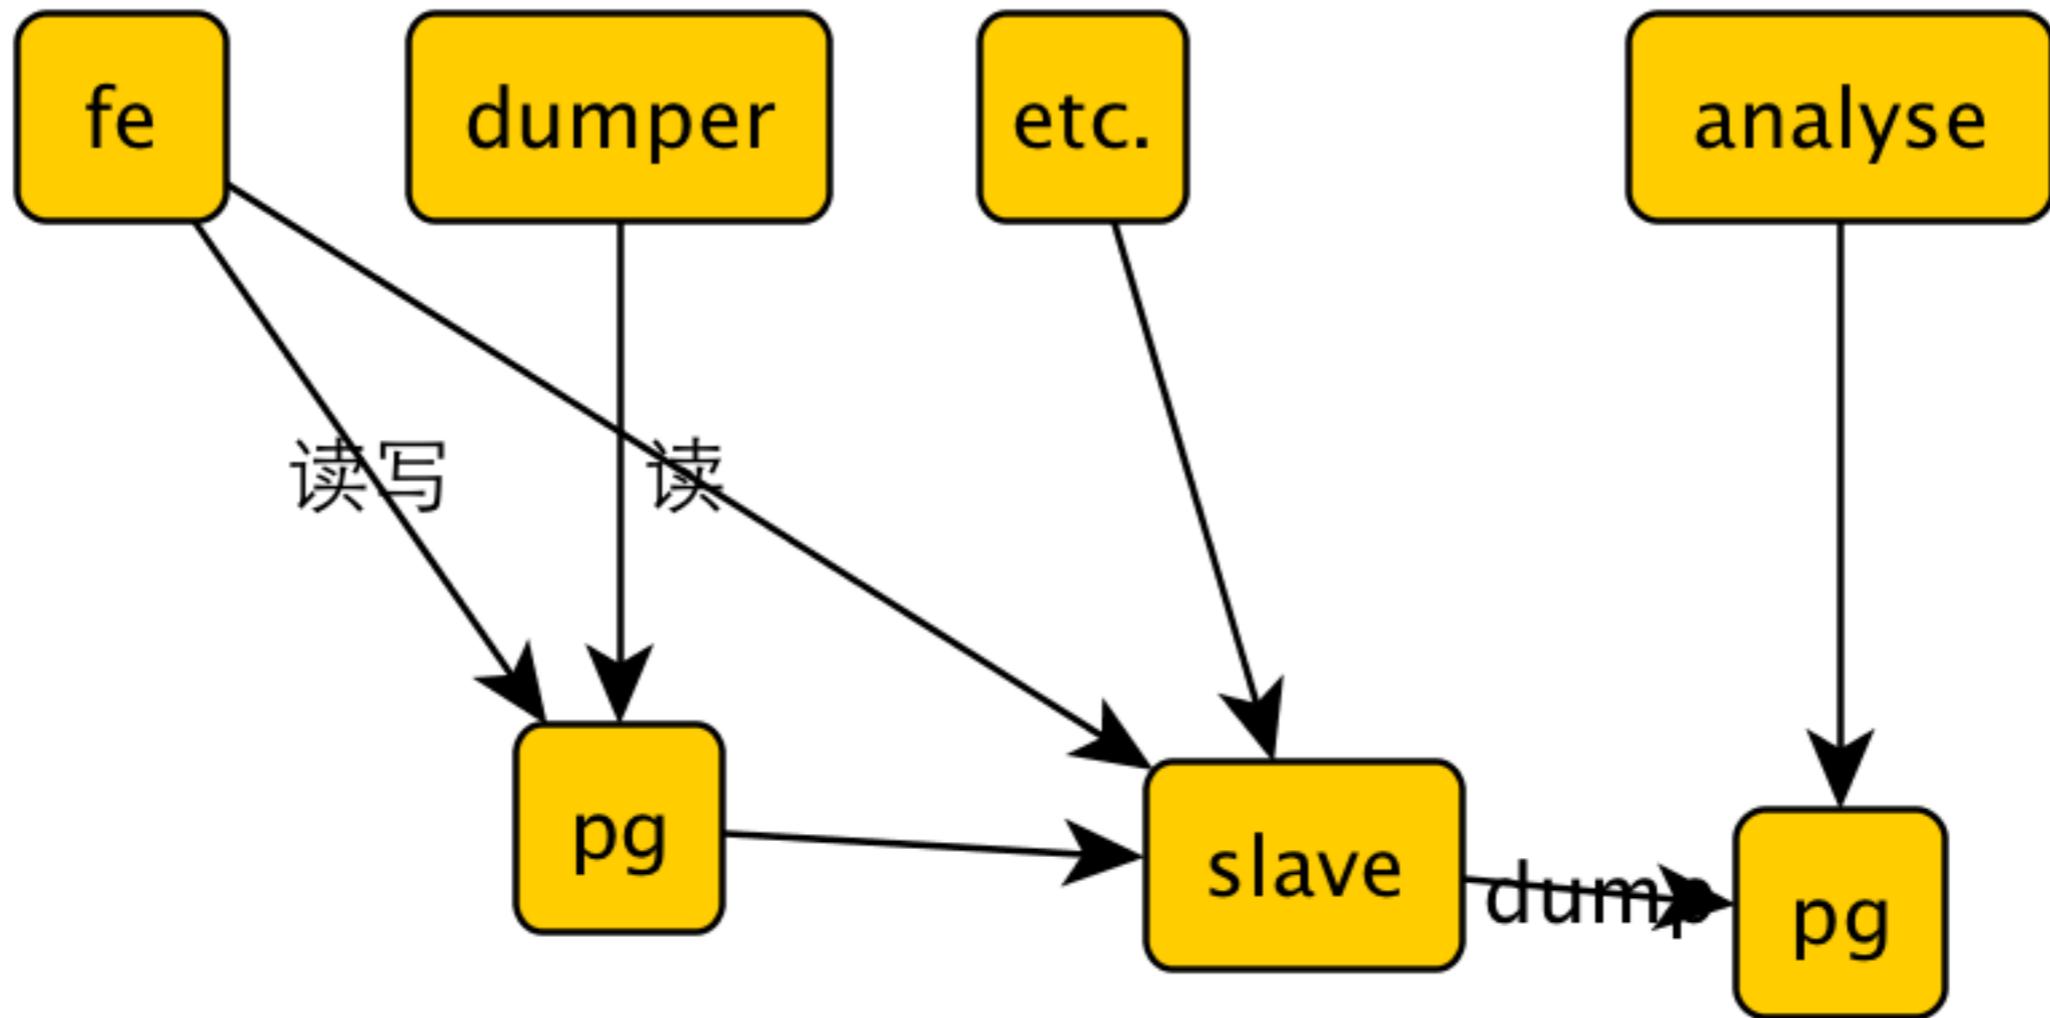
# 改进



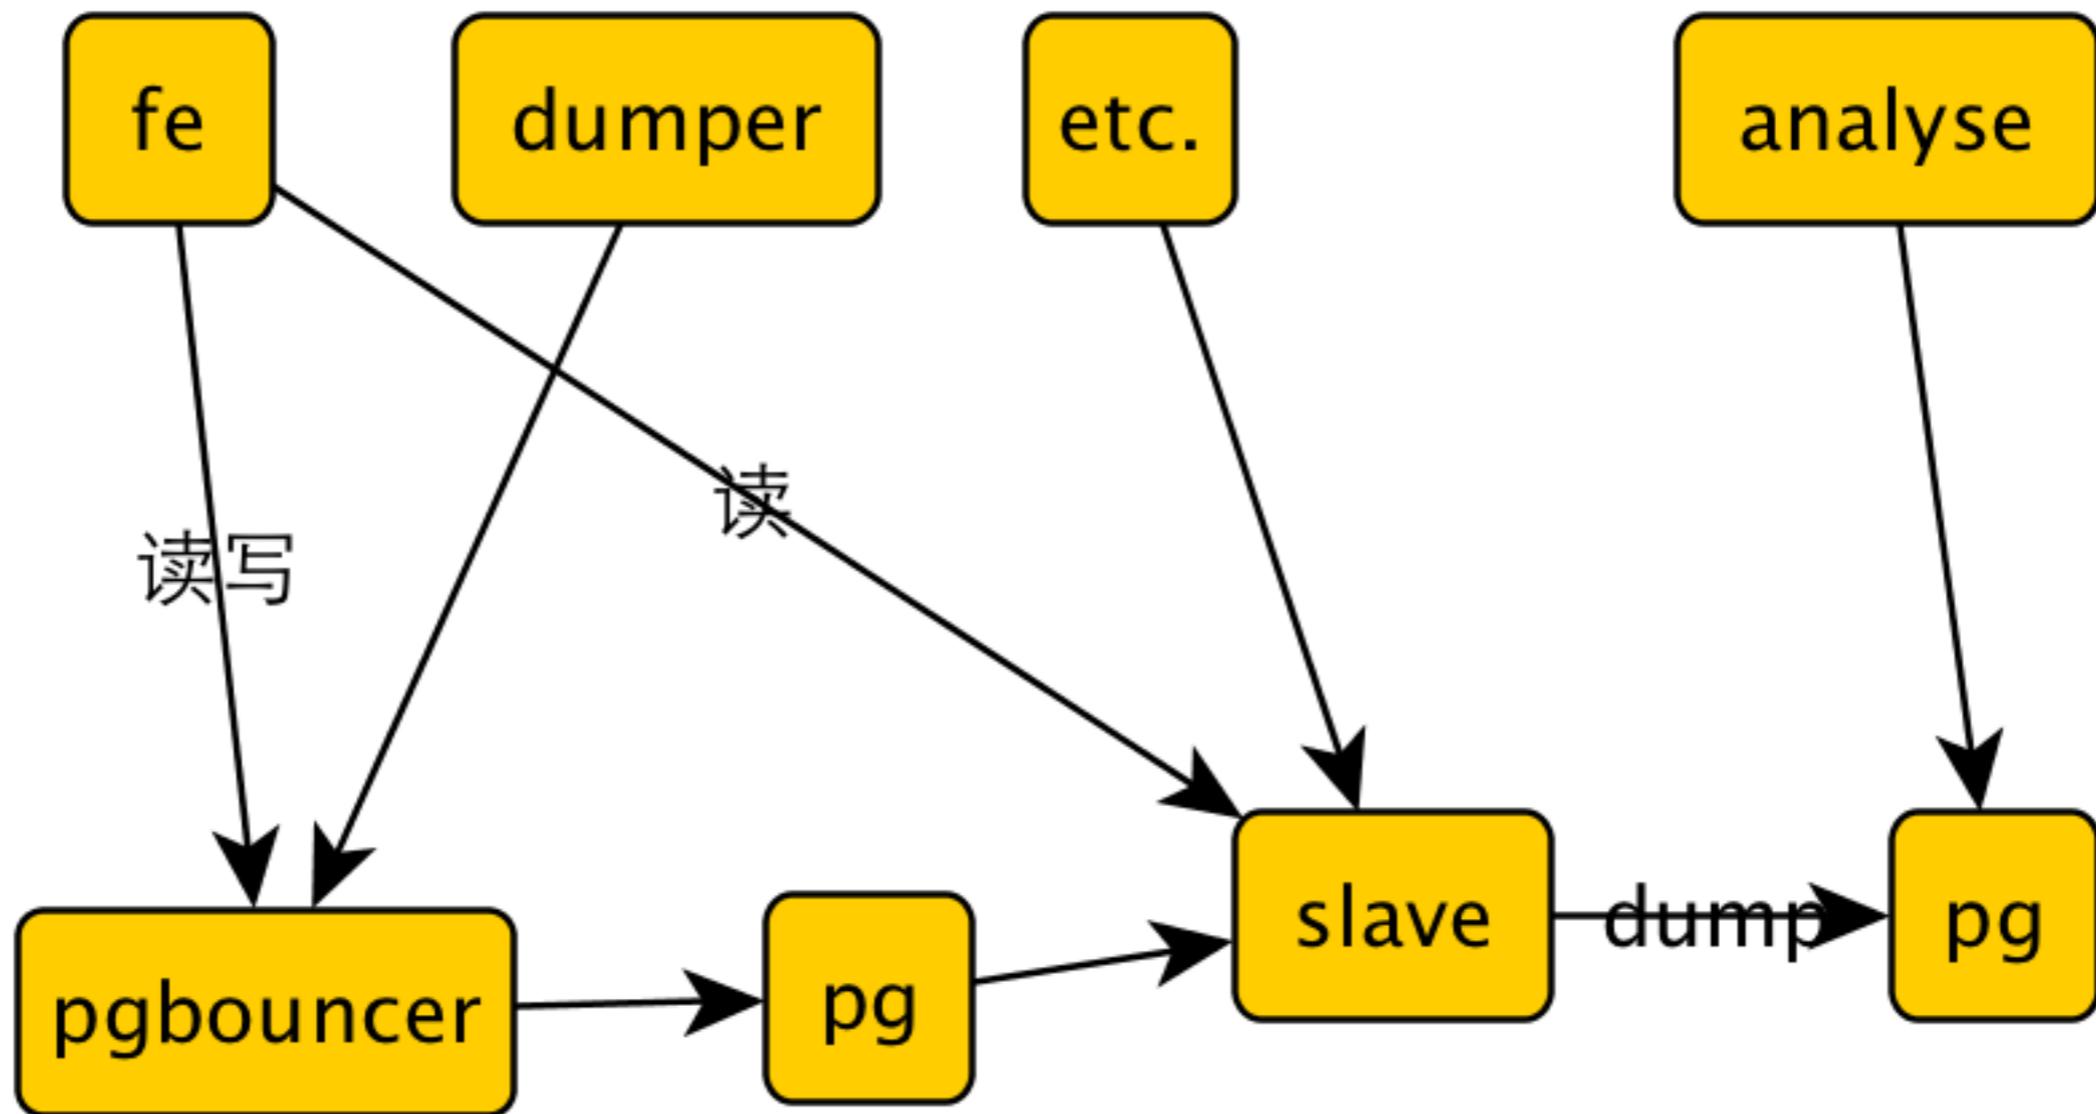
# 改进



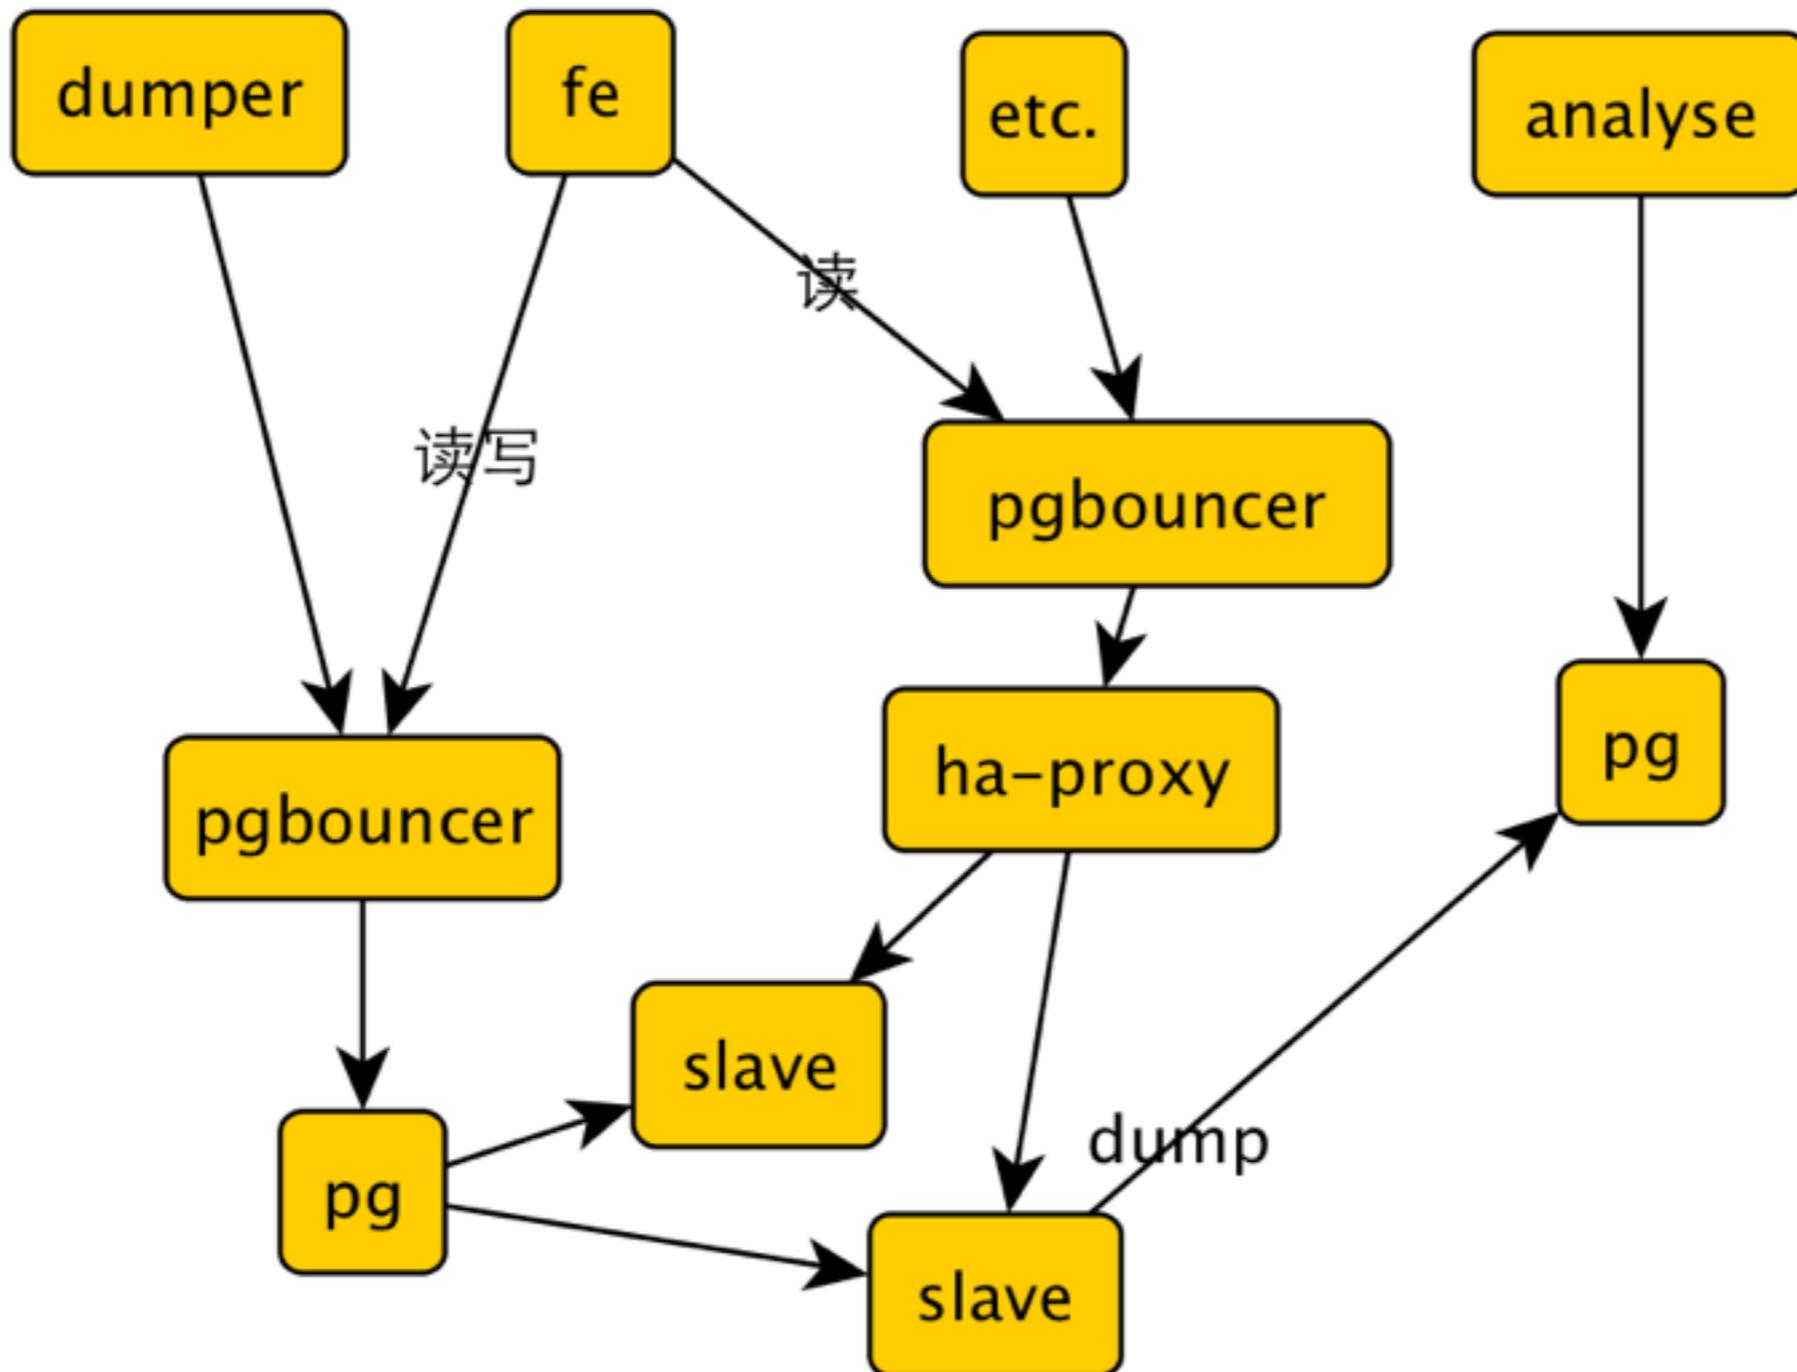
# 第二版结构

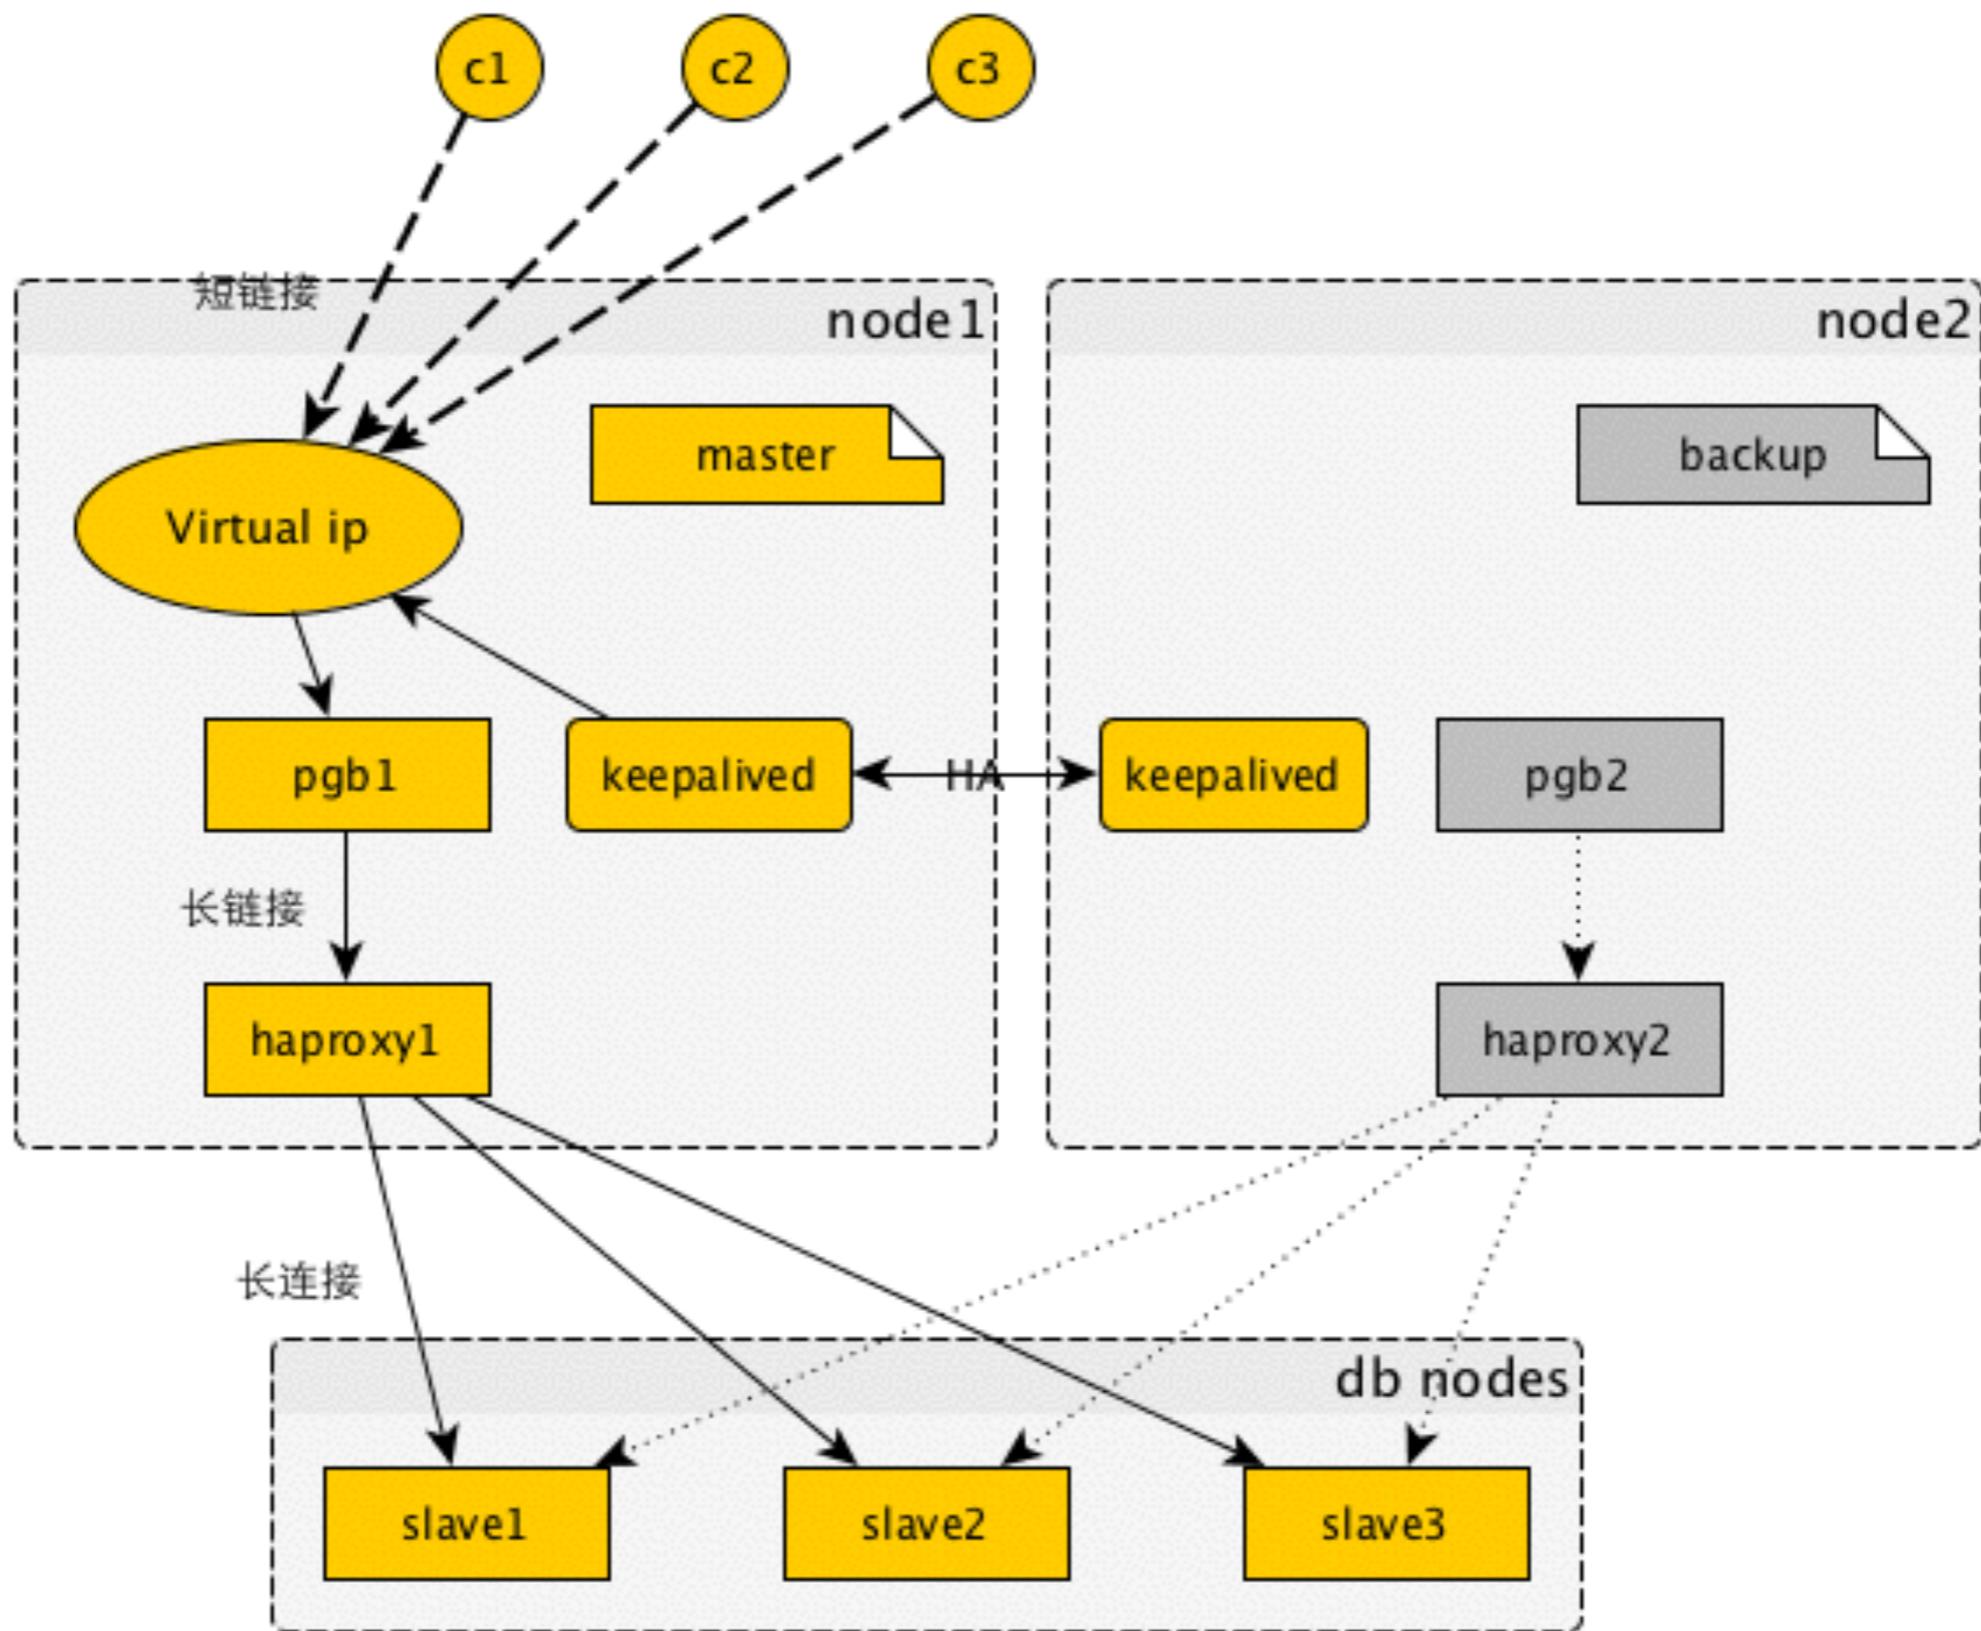


# 改进



# 改进





# 改进

- pgbouncer 补丁
  - ip 限制
  - ip 透传
- [http://qunardba.gitcafe.io/2015/04/03/pgbouncer-ip\\_limit/](http://qunardba.gitcafe.io/2015/04/03/pgbouncer-ip_limit/)

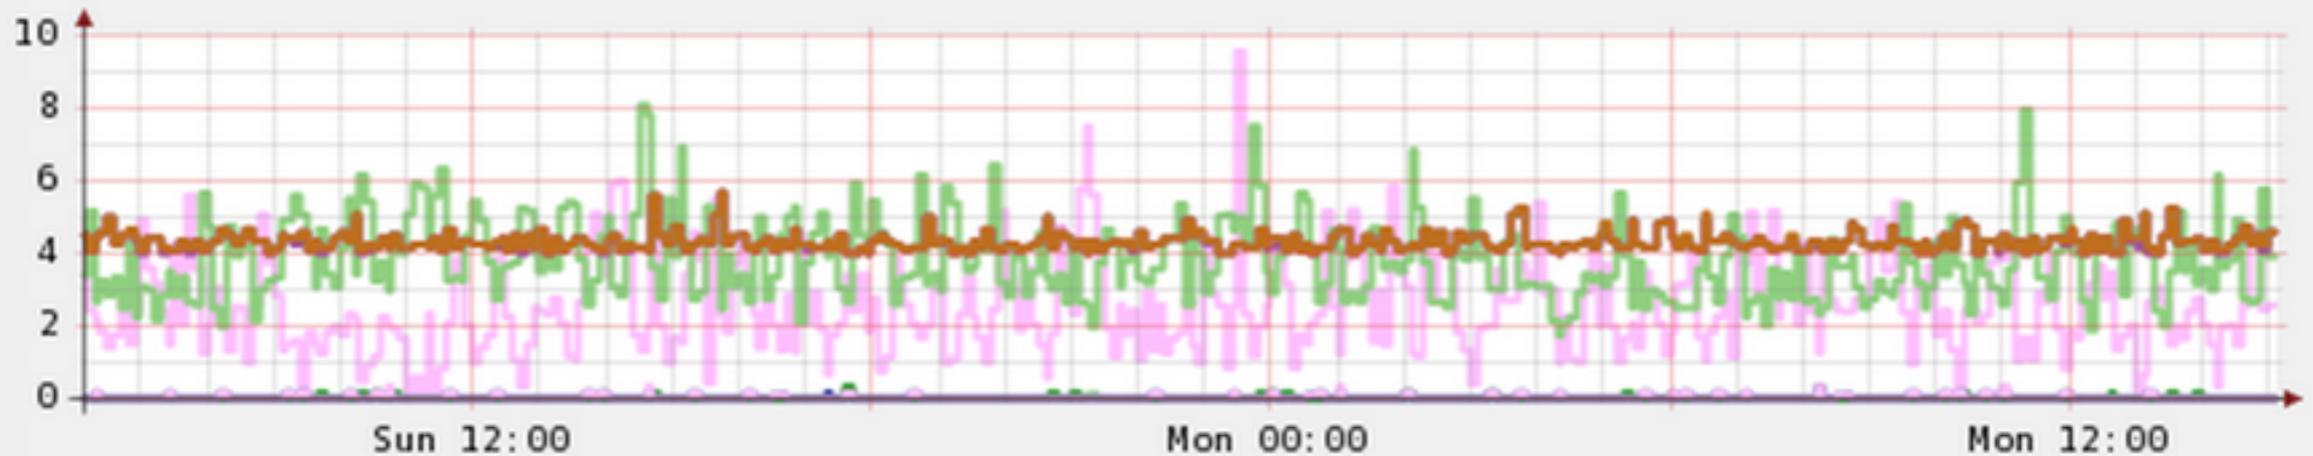
# 总结

- 基于 Master-slave 结构存在的问题
  - 同步延迟
  - Conflict with recovery
  - 遇到数据损坏集群一次全部 down 机
- 大版本升级需要停机，从库需要重做
- 通过 pgbouncer 可以有效的控制 pg 的连接数

# 监控措施

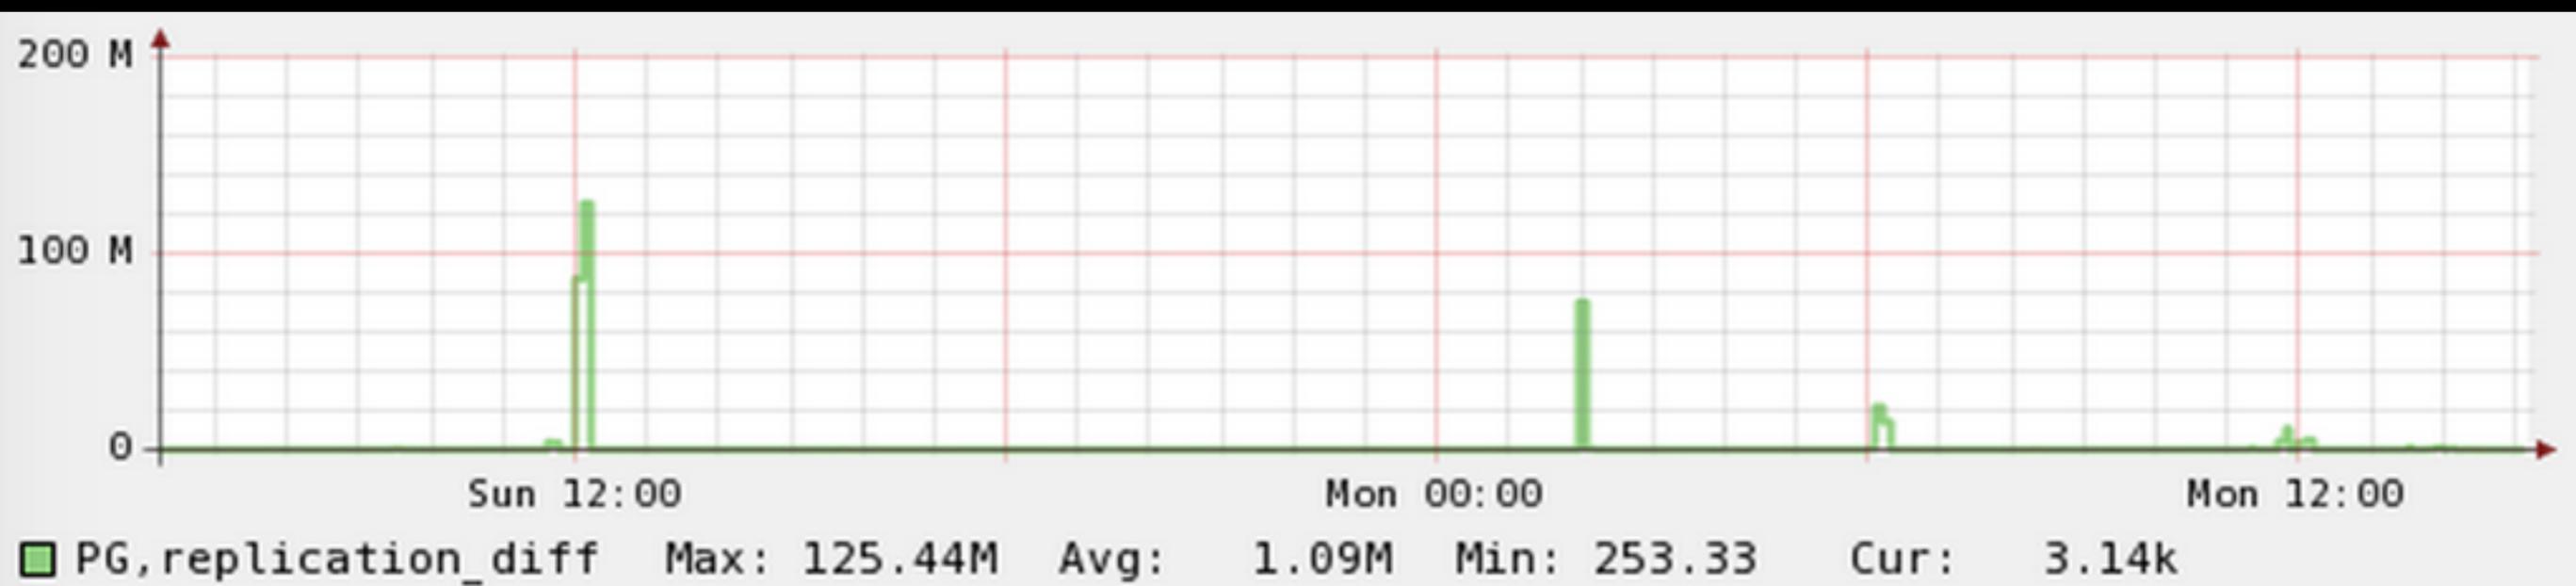
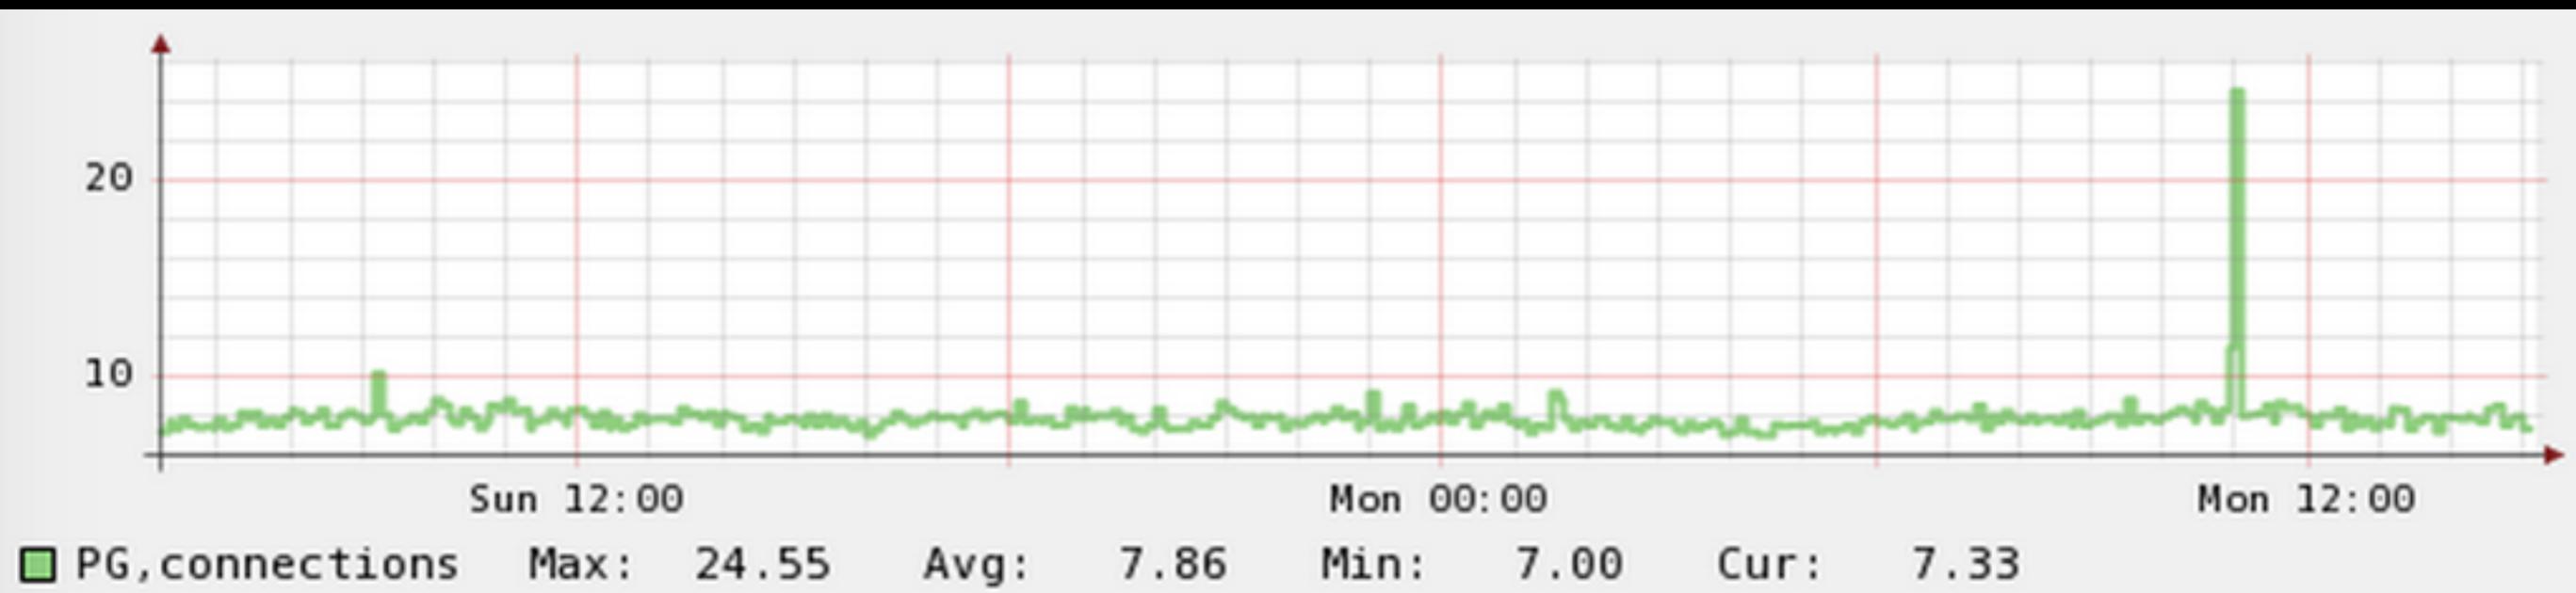
- 监控是运维的核心工作
- 好的运维需要能在问题出现前发现/避免问题

# pgbouncer

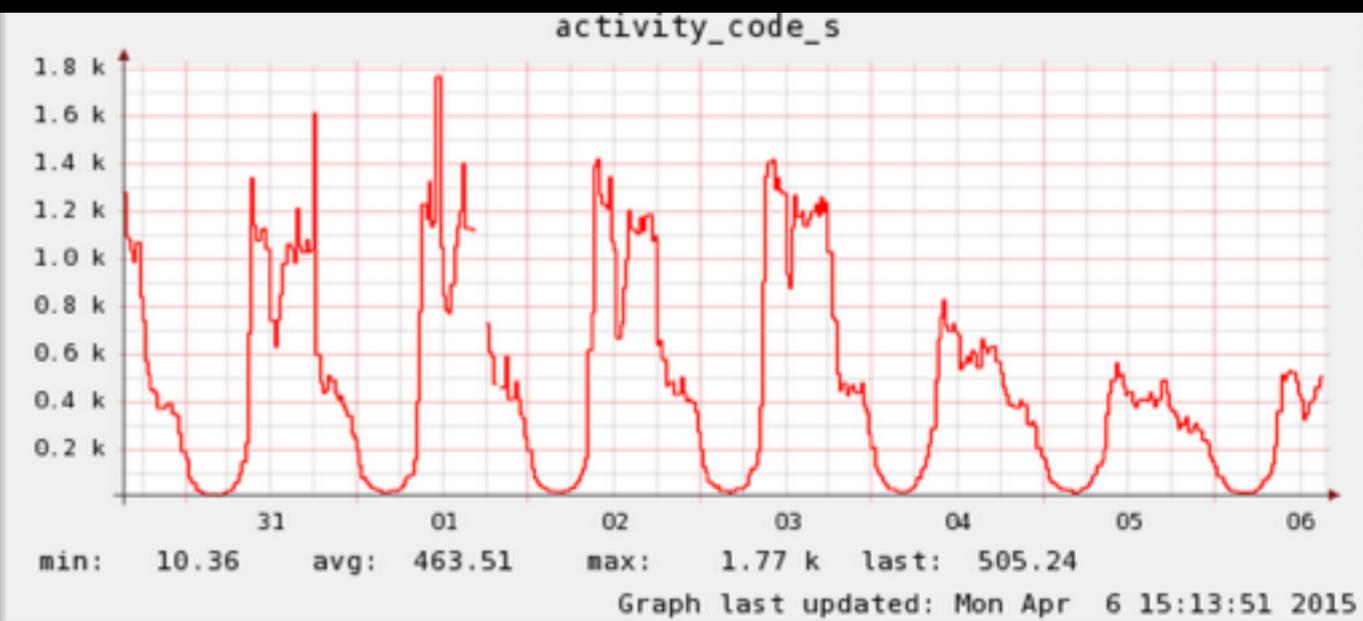
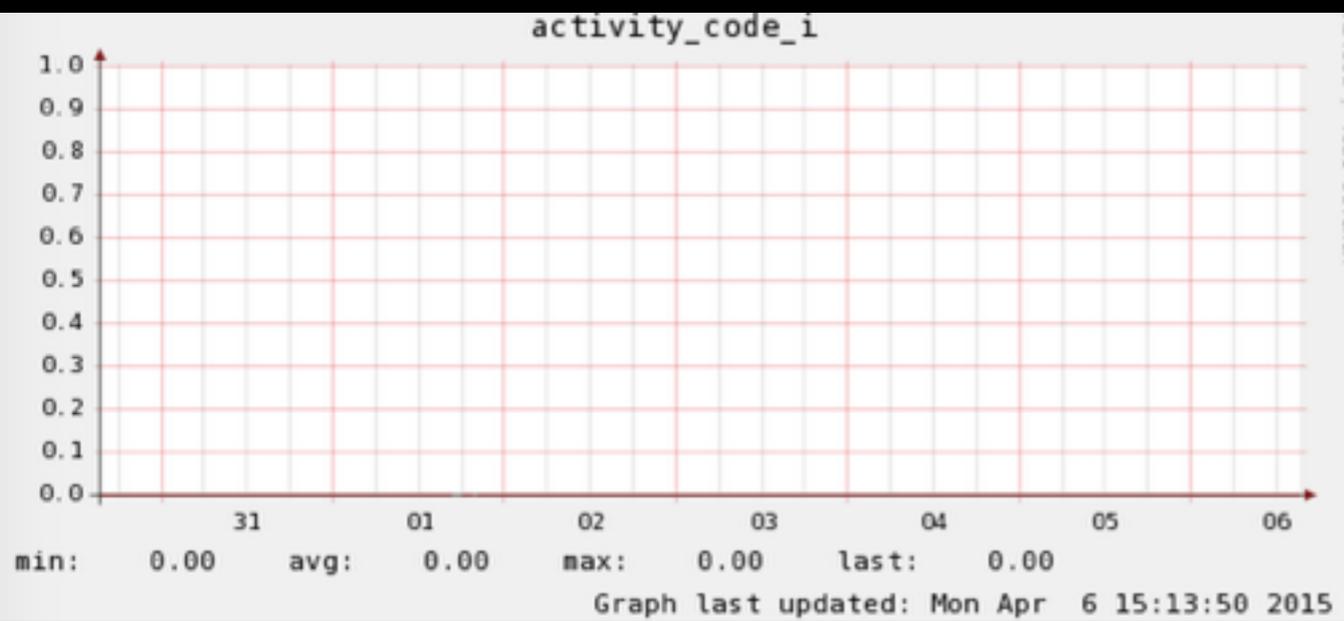


vacationd_sv_used, vacationd_sv_used	Max:	nan	Avg:	nan	Min:	nan	Cur:
vacationd_sv_login, vacationd_sv_login	Max:	nan	Avg:	nan	Min:	nan	Cur:
vacation_maxwait, vacation_maxwait	Max:	0.00	Avg:	0.00	Min:	0.00	Cur:
vacationd_sv_idle, vacationd_sv_idle	Max:	nan	Avg:	nan	Min:	nan	Cur:
vacation_sv_used, vacation_sv_used	Max:	9.55	Avg:	2.54	Min:	166.67m	Cur:
vacationd_sv_active, vacationd_sv_active	Max:	nan	Avg:	nan	Min:	nan	Cur:
vacationd_cl_waiting, vacationd_cl_waiting	Max:	nan	Avg:	nan	Min:	nan	Cur:
vacationd_maxwait, vacationd_maxwait	Max:	nan	Avg:	nan	Min:	nan	Cur:
vacation_sv_idle, vacation_sv_idle	Max:	8.09	Avg:	3.81	Min:	1.74	Cur:
vacation_sv_tested, vacation_sv_tested	Max:	333.33m	Avg:	10.10m	Min:	0.00	Cur:
vacationd_cl_active, vacationd_cl_active	Max:	nan	Avg:	nan	Min:	nan	Cur:
vacation_sv_login, vacation_sv_login	Max:	333.33m	Avg:	22.22m	Min:	0.00	Cur:
vacation_cl_waiting, vacation_cl_waiting	Max:	333.33m	Avg:	23.74m	Min:	0.00	Cur:
vacation_sv_active, vacation_sv_active	Max:	5.68	Avg:	4.32	Min:	4.00	Cur:
vacationd_sv_tested, vacationd_sv_tested	Max:	nan	Avg:	nan	Min:	nan	Cur:
vacation_cl_active, vacation_cl_active	Max:	5.68	Avg:	4.33	Min:	4.00	Cur:

# pg



# pg



# pg

**db\_stats** 15 minutes ago to a few seconds ago refreshed every 1m

**QUERY**

Query Generator

**FILTERING**

time must  
field : @timestamp  
from : now-15m  
to : now

**DATABASES**

b2c\_product (13110) Missing field (0) Other values (0)

**CALLS**

calls total per 10s (13110 hits)

**TOTAL ROWS**

**DETAIL**

0 to 100 of 500 available for paging

Fields: @timestamp, @version, \_id, \_index

calls	rows	query	md5_query
140	6		bd0856fd065b19c91dcbfa8fe567fef
23	21		f6e8b2c69768be3e9f9aa97a086f1635
2334	2334		2db7e21284514393a0d1bac39ac9e516
111	111		18c8c668d5c16c8689bea75762873ddc
93	7		c8ae77580599ac8cfd5e60df31449059

# pg

- locks
- slow\_query
- <http://qunardba.gitcafe.io/2015/01/18/maintenance-tips/>

# 将来的改进

- BDR
- PXC
- 基于 pgbouncer 分离监控层

# 广告时间

- <http://qunardba.gitcafe.io/>



A screenshot of the 'Qunar PostgreSQL DBA Blog' website. The header is orange with a blue ghost icon and the text 'QUNAR PostgreSQL DBA Blog' and 'All About PostgreSQL'. Navigation links for 'Home' and 'Archives' are on the right. The main content area lists several articles: 'PostgreSQL BDR testing' (dated Apr 18 2015), 'pgbouncer ip 限制和透传 ip 补丁' (dated Apr 3 2015), 'pg\_hint\_plan' (dated Feb 16 2015), and 'pg94 不同 blocksize 性能比较'. A 'Tags' sidebar on the right includes 'BDR 1', 'explai', and 'logs'.

# 广告时间

- 招聘
  - PostgreSQL DBA
  - 数据统计分析工程师
  - Java/IOS 开发工程师
  - [dong.wang@qunar.com](mailto:dong.wang@qunar.com)

Q&A