



SCHIBSTED
CLASSIFIED MEDIA

2TB of RAM ought to be enough for anybody

PG Day FR 2014

The Presenter



Jean-Louis Bergamo - @jlb666
Paris, France

The Presenter – short bio



- 2000-2006 : Easynet
- 2007-2010 : Skyrock.com
 - European biggest blog platform
- 2010-2011 : Winamax.fr
 - #1 Online poker game
- 2011-2012 : Cadremploi.fr
- 2013-03-01: Infra manager



leboncoin.fr
vendez, achetez, près de chez vous



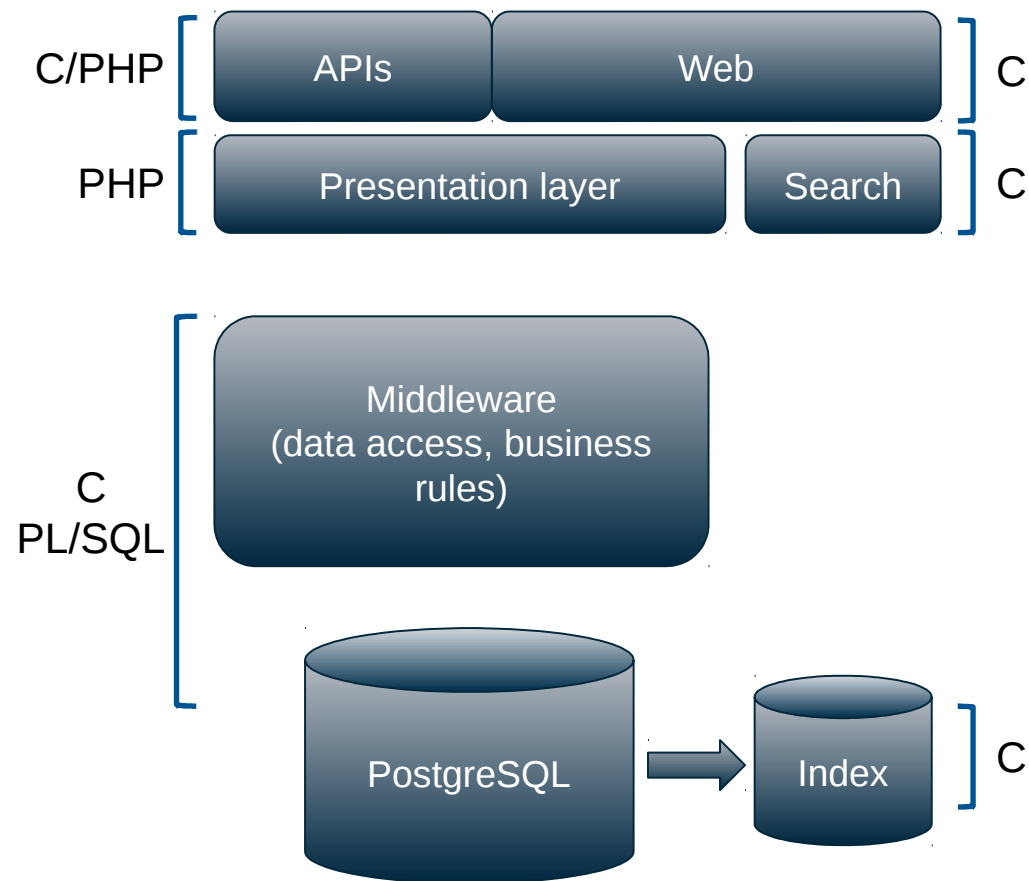
SCHIBSTED
CLASSIFIED MEDIA

Schibsted Classified Media

From 1 to 30+ Countries in 7 years



Templated deployment in 30 countries w/ shared technology



- Technology originally inherited from Blocket.se
- Has since evolved to power all 30 sites
- Focus on performance and ease of local modifications

PostgreSQL in the SCM Platform

100+ servers running PostgreSQL

8TB of data

50+ million classified ads



Schibsted Classified Media & PostgreSQL married...and in love





SCHIBSTED
CLASSIFIED MEDIA



#1 Classified Web site in France

History

Project initiated in 2006

Site launch: early 2007

Based on **blocket** technology


from 1 to 230 people, challenger to #1 in 7 years



leboncoin.fr
vendez, achetez, près de chez vous

Leboncoin.fr part d'une idée simple : la bonne affaire est au coin de la rue ! Pour passer ou chercher des annonces, cliquez sur la région de votre choix et trouvez la bonne affaire parmi **24 612 361 annonces**.

Simple, rapide et efficace !



Alsace
Aquitaine
Auvergne
Basse-Normandie
Bourgogne
Bretagne
Centre
Champagne-Ardenne
Corse
Franche-Comté
Haute-Normandie
Ile-de-France
Languedoc-Roussillon
Limousin
Lorraine
Midi-Pyrénées
Nord-Pas-de-Calais
Pays de la Loire
Picardie
Poitou-Charentes
Provence-Alpes-Côte d'Azur
Rhône-Alpes

Guadeloupe
Martinique
Guyane
Réunion

[ACCUEIL](#)

[DÉPOSER UNE ANNONCE](#)

[OFFRES](#)

[DEMANDES](#)

[MES ANNONCES](#)

[BOUTIQUES](#)

[MON COMPTE](#)

[AIDE](#)

Toutes catégories

Ile-de-France

Ville ou code postal

☐ Recherche dans le titre uniquement

☐ Annonces ★ Urgentes uniquement

CHERCHER

» Comment effectuer une recherche par mots-clés ?

♥ Consultez des annonces de célibataires. Inscription gratuite !

TOUTES

1 - 35 de **3 713 152**

PARTICULIERS

3 392 246 annonces

PROFESSIONNELS

320 906 annonces

[Cacher les photos](#) | [Trier par prix](#)

Hier
23:19



3

Sac à main en cuire bleu foncé

Accessoires & Bagagerie
Paris 11ème

80 €

Hier
23:19



3

Apple iMac Blanc

Informatique
Eaubonne / Val-d'Oise

350 €

Hier
23:19



2

Tee shirt Gap taille S

Vêtements
Paris 13ème

6 €

ANNONCES A LA UNE



**APPARTEMENT 3 PIECES
65m2 LE CHESNAY**

328 000 €




Ligne complète racing

Accueil > Val-d'Oise > Informatique > Apple iMac Blanc

Apple iMac Blanc

Mise en ligne par **Pépin3pommes** le 17 mars à 23:19.



3 PHOTOS DISPONIBLES

Prix: 350 €

Ville : **Eaubonne**

Code postal : **95600**

€ [Simulez votre crédit](#)

[Votre box B&YOU pour 15.99€/mois](#)

Description :

Vends Apple iMac Blanc
Ecran 20 pouces (1680 x 1050 pixels)

Contactez l'annonceur

[Envoyer un email](#)

[Voir le numéro de téléphone](#)

(Je refuse tout démarchage commercial)

Gérer votre annonce

[Modifier](#)

[Supprimer](#)

[Remonter en tête de liste](#)

[Mettre en avant](#)

Simulez votre crédit gratuitement

Etudiez votre financement : calculez vos mensualités en ligne. Réponse de principe immédiate.

[> Cliquez ici](#)

PROMO BOX INTERNET
-50% SUR L'ACHAT DE VOTRE BOX **B.YOU**
JUSQU'AU 3 MARS SEULEMENT !

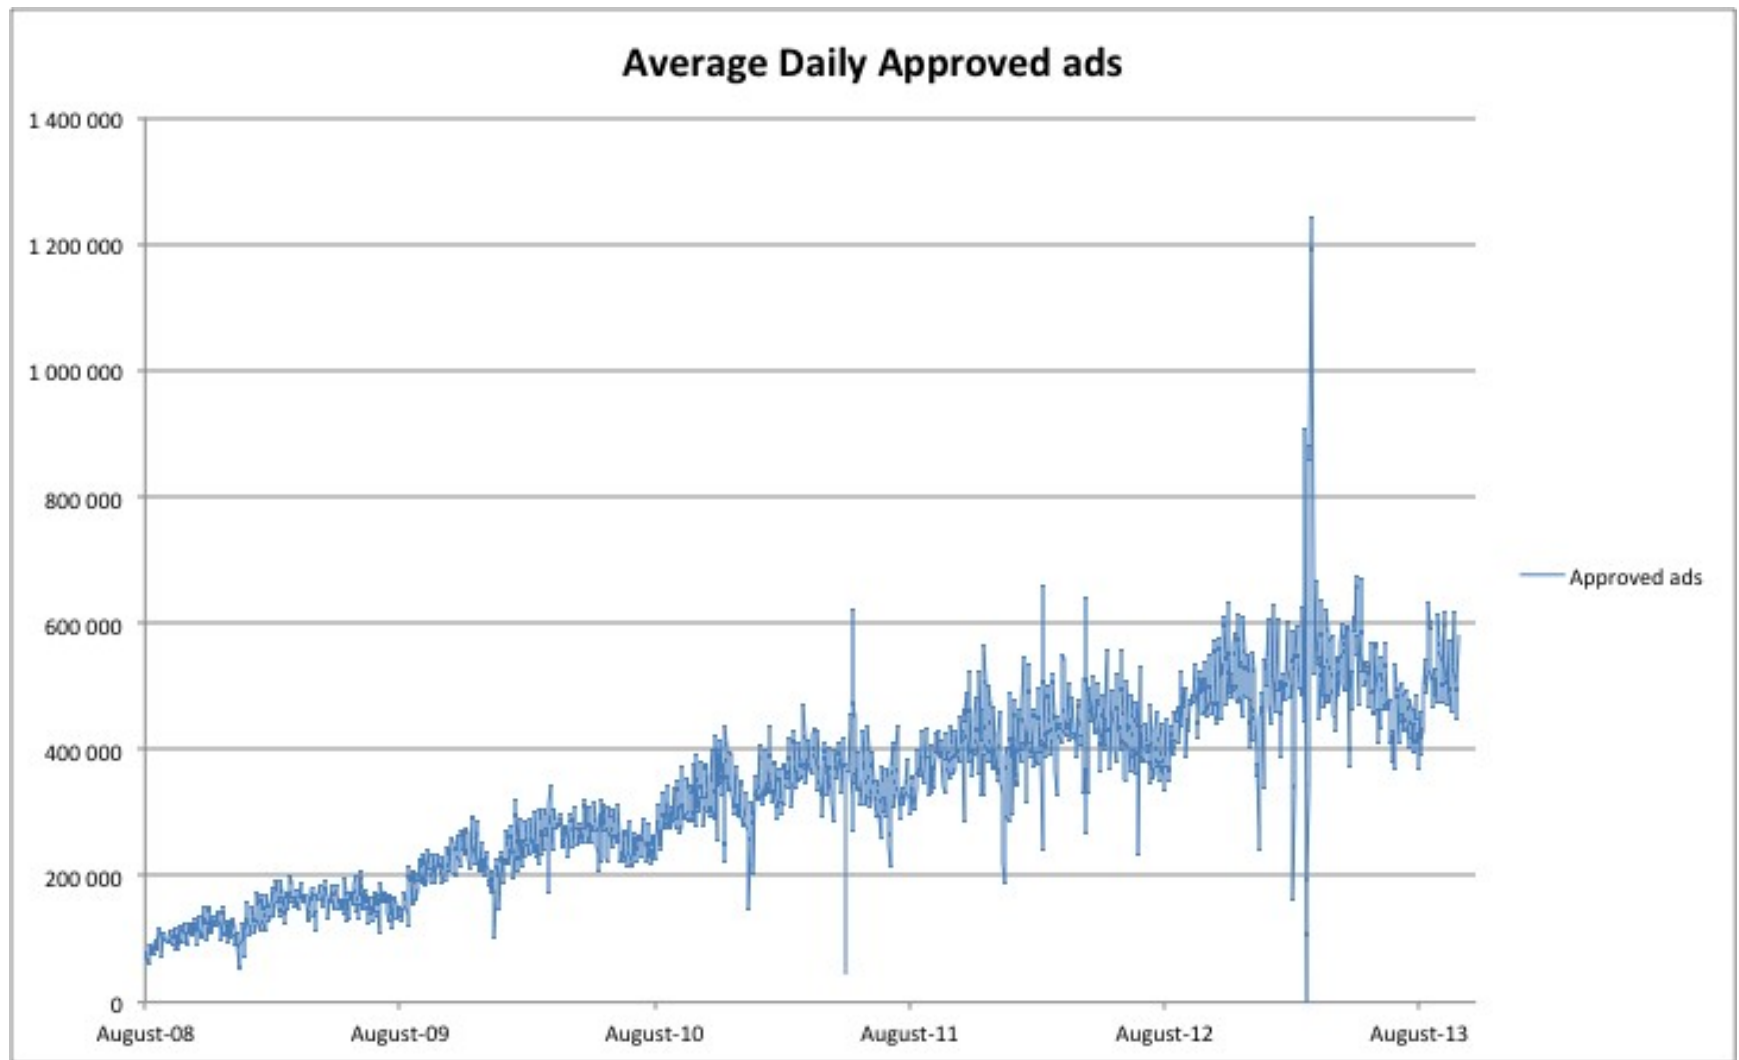
Mardi 13 mars

minimum **15** MILLIONS D'EUROS à gagner

* Montant à partager au rang 1. Voir règlement.

Jouer comporte des risques : isolement, endettement... Appelez le 09-74-75-13-13 (appel non surtaxé)

Explosive growth...with a few bumps along the way!



Big Audience

250M page views / day

5M unique visitors / day

18M UV / month

That's 1/3 of French internet population...

600000+ new ads / day

25M live ads

#7 most visited Website in France

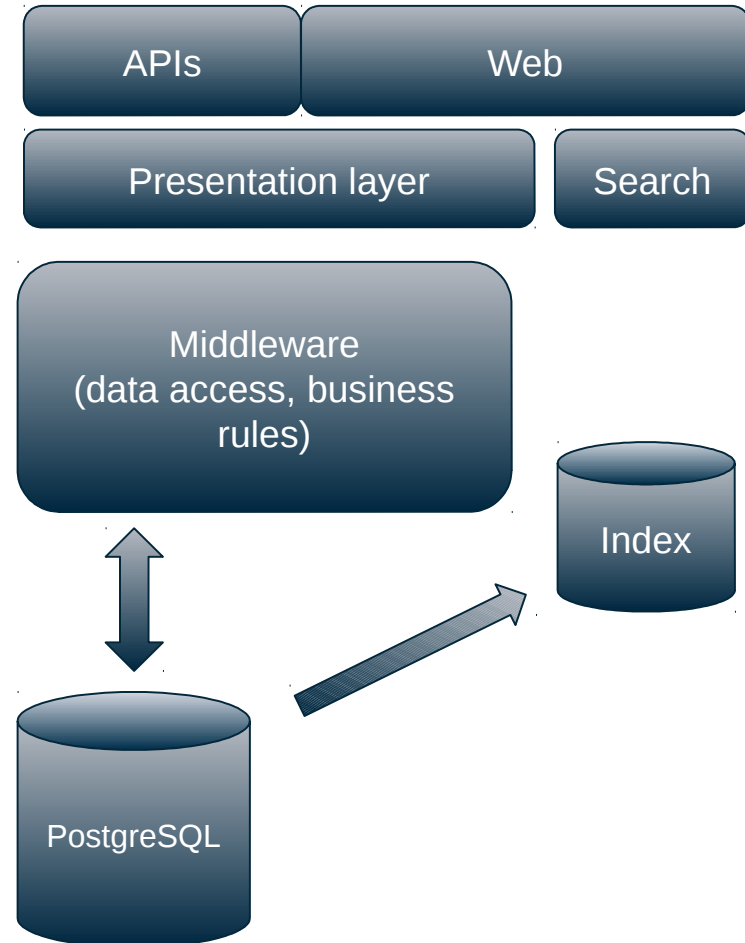
Big Ops

300+ servers in 2 DCs

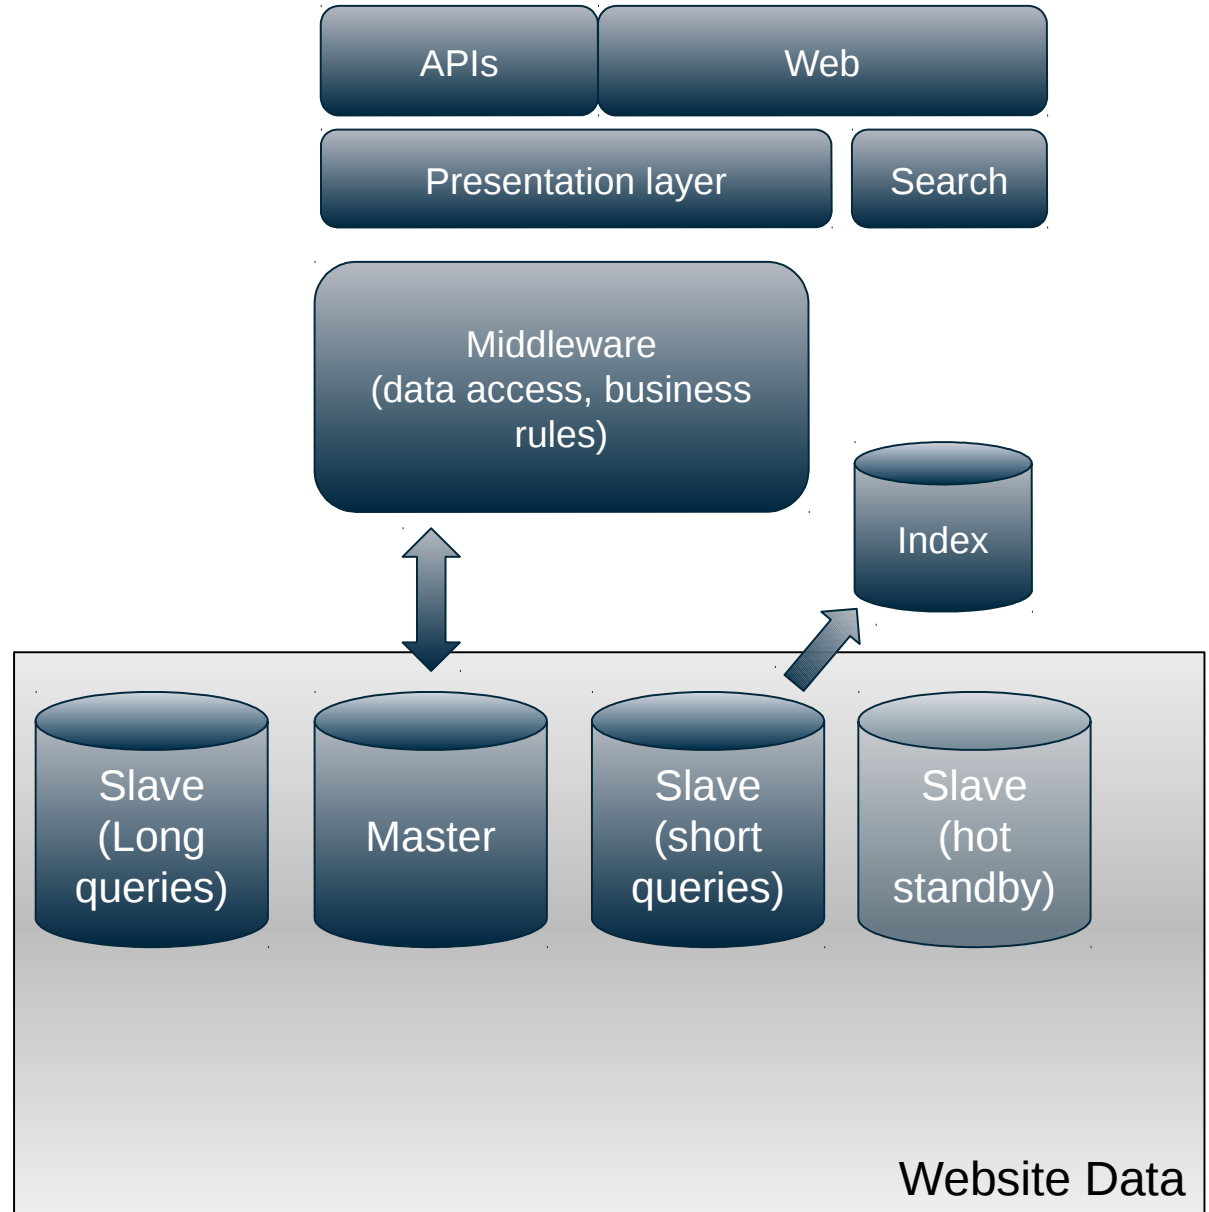
20 servers hosting PG databases
(in production)



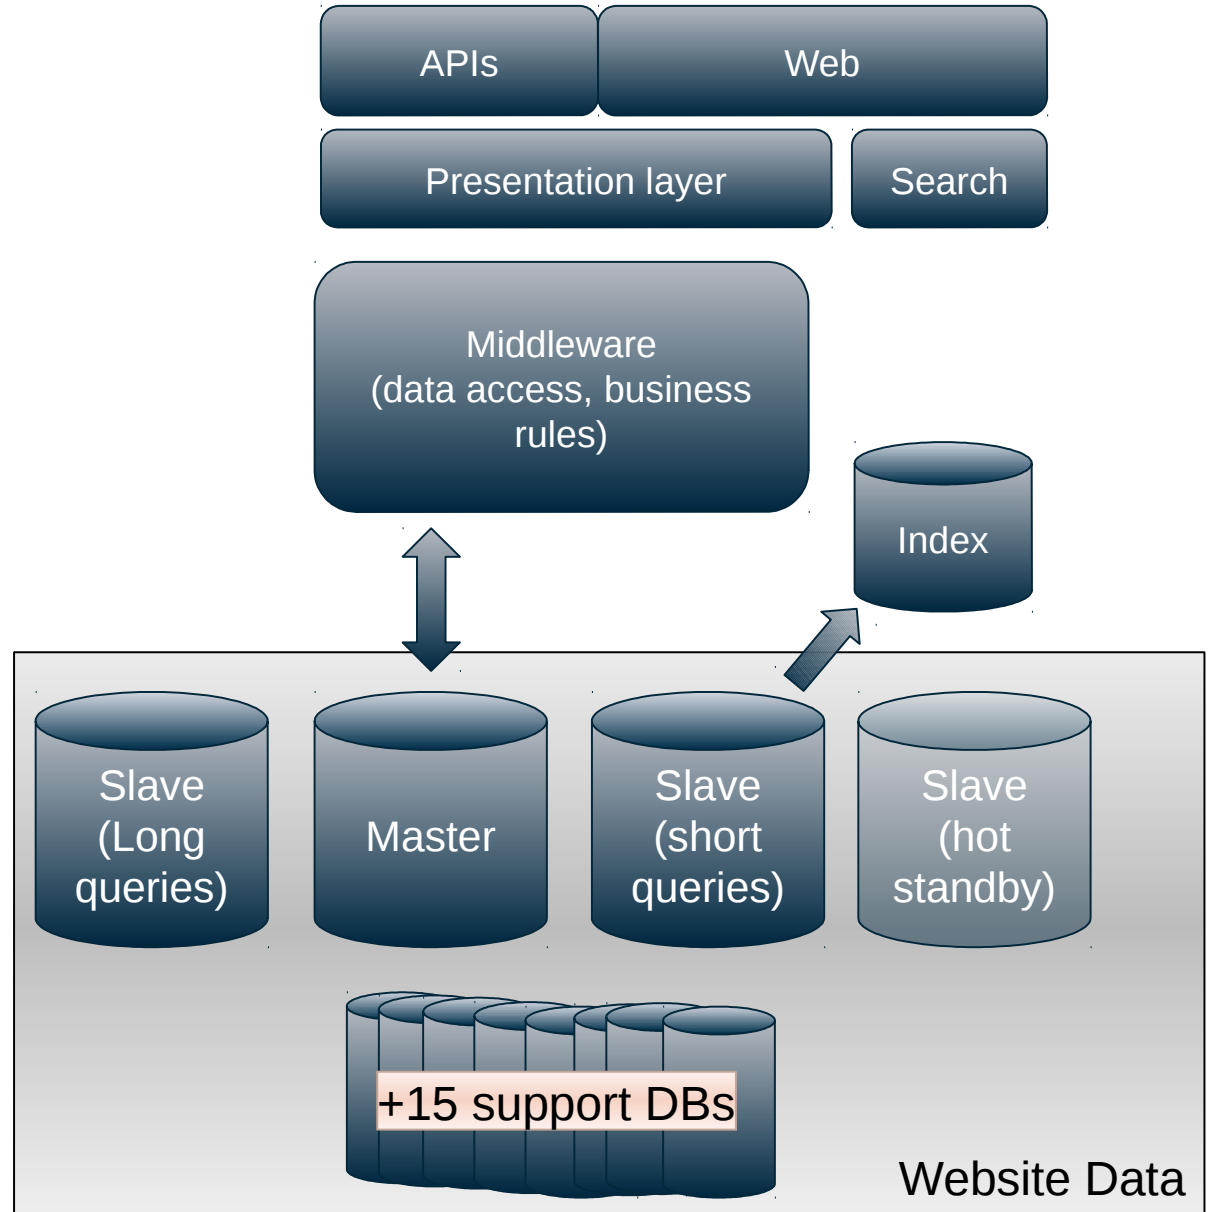
Built on SCM Technology



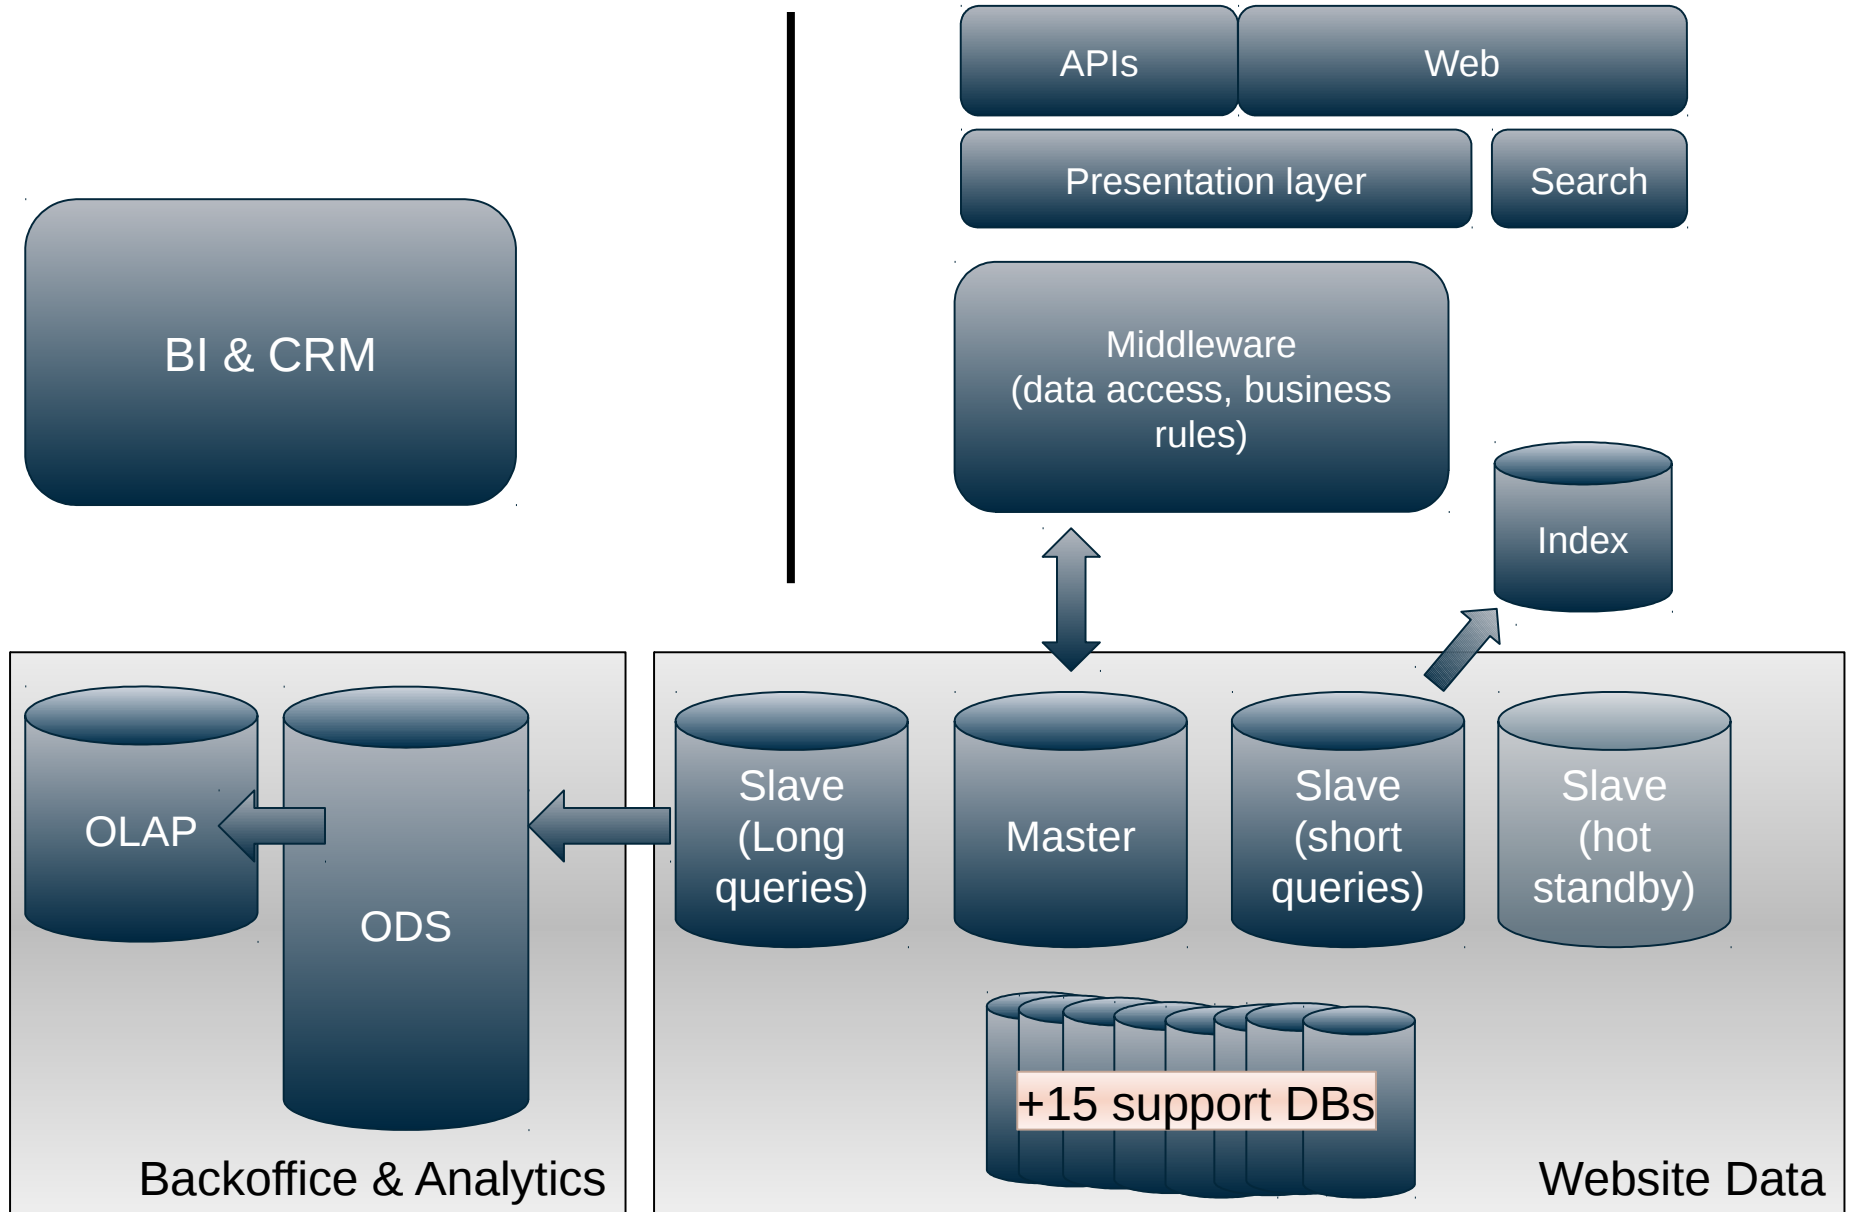
Built on SCM Technology



Built on SCM Technology

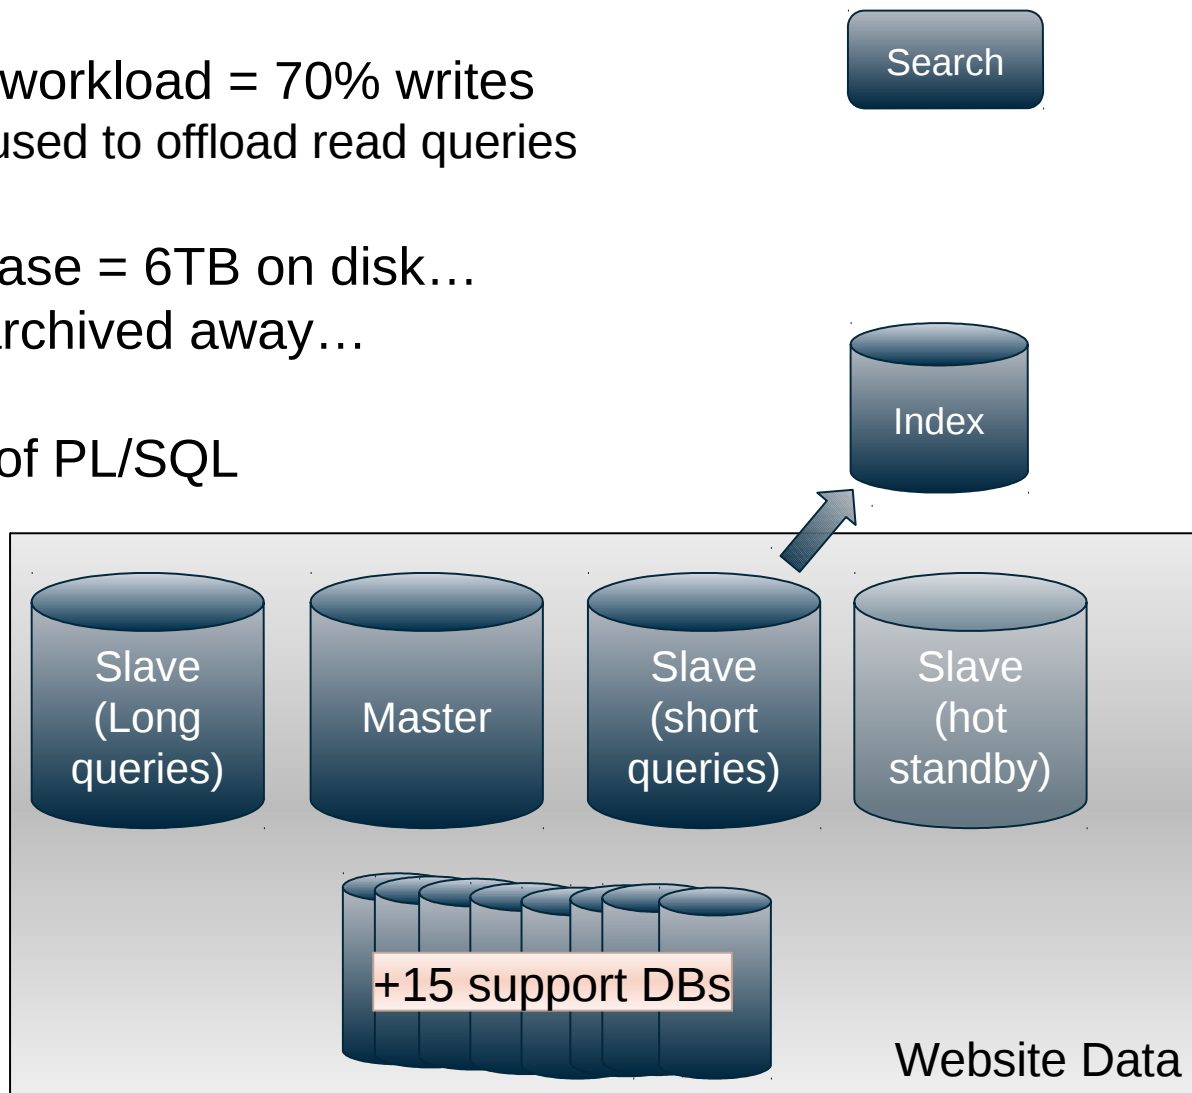


We use PostgreSQL everywhere!

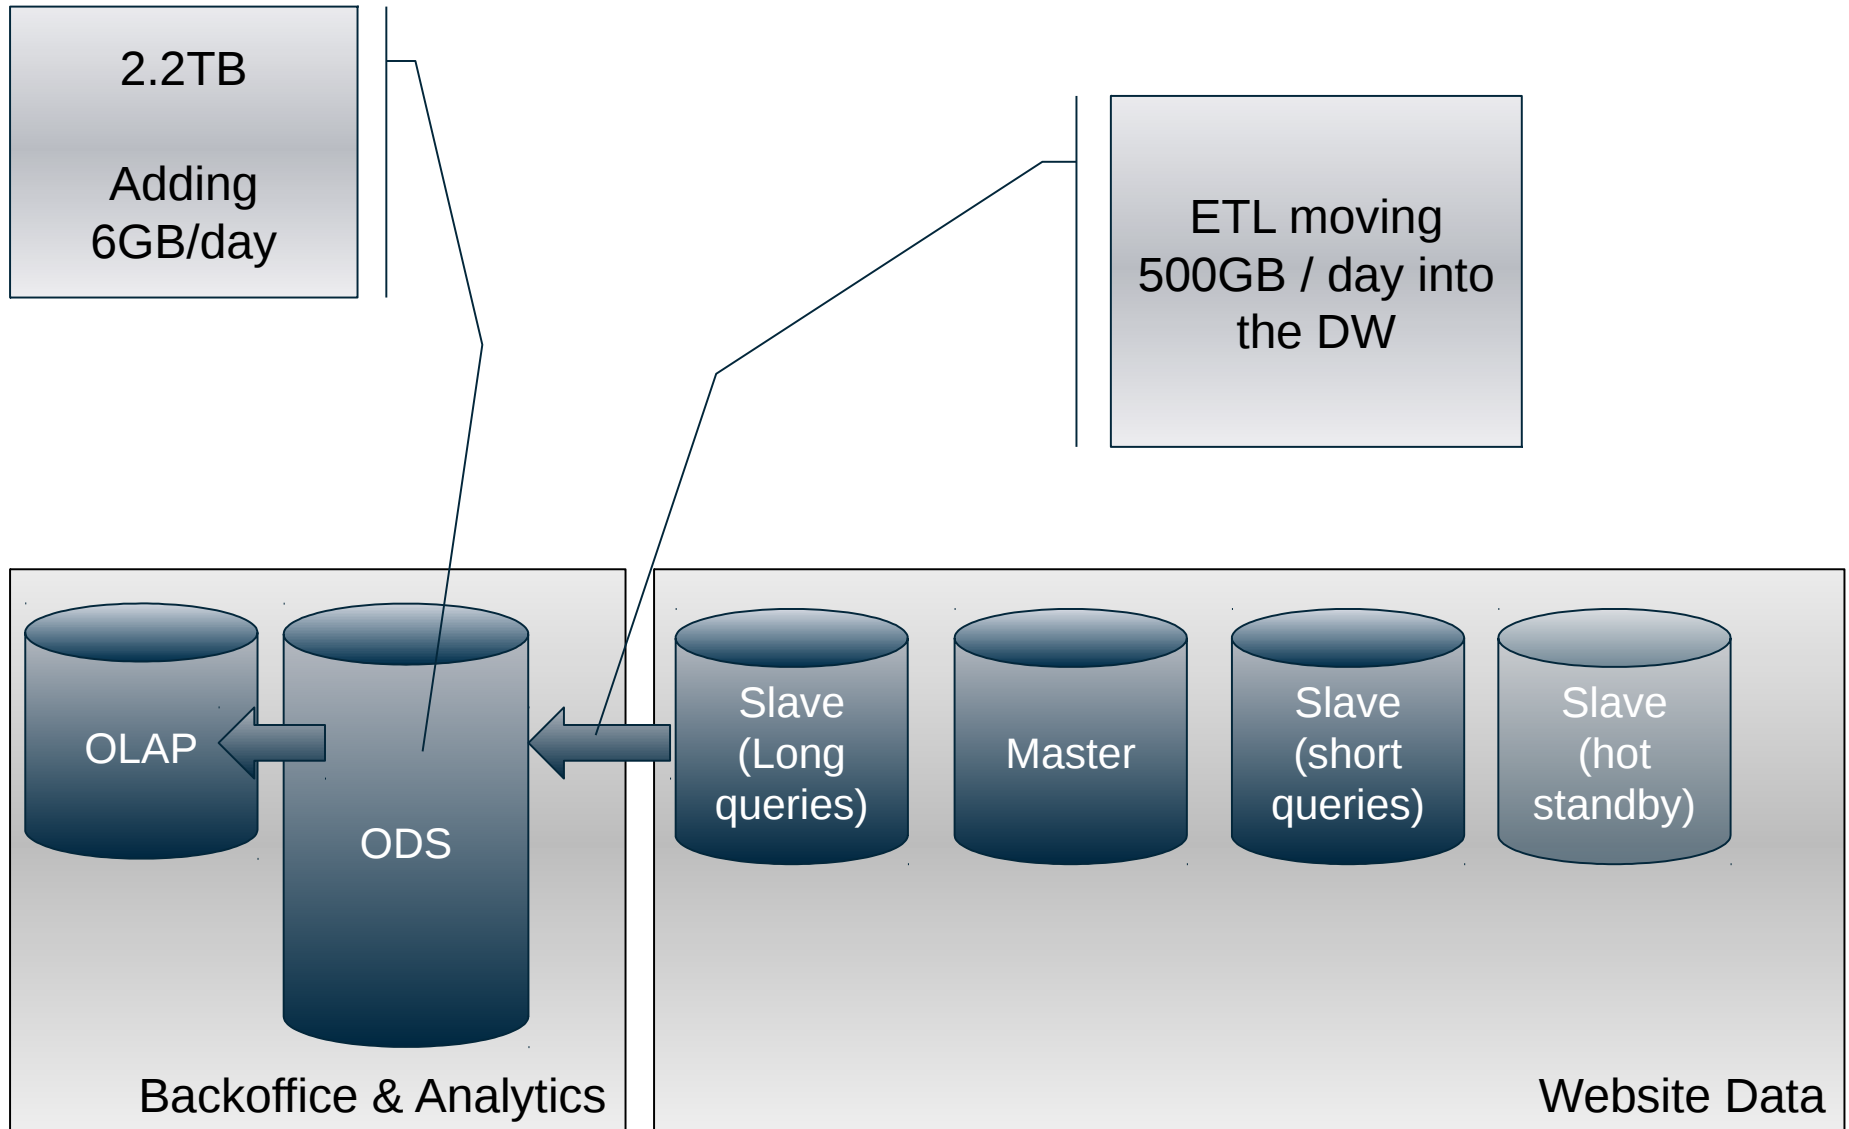


We try to limit writes!


- Index/Search acting as a structured cache
- Master DB workload = 70% writes
 - Slaves used to offload read queries
- Main database = 6TB on disk...
 - +4TB archived away...
- 20K LOCs of PL/SQL



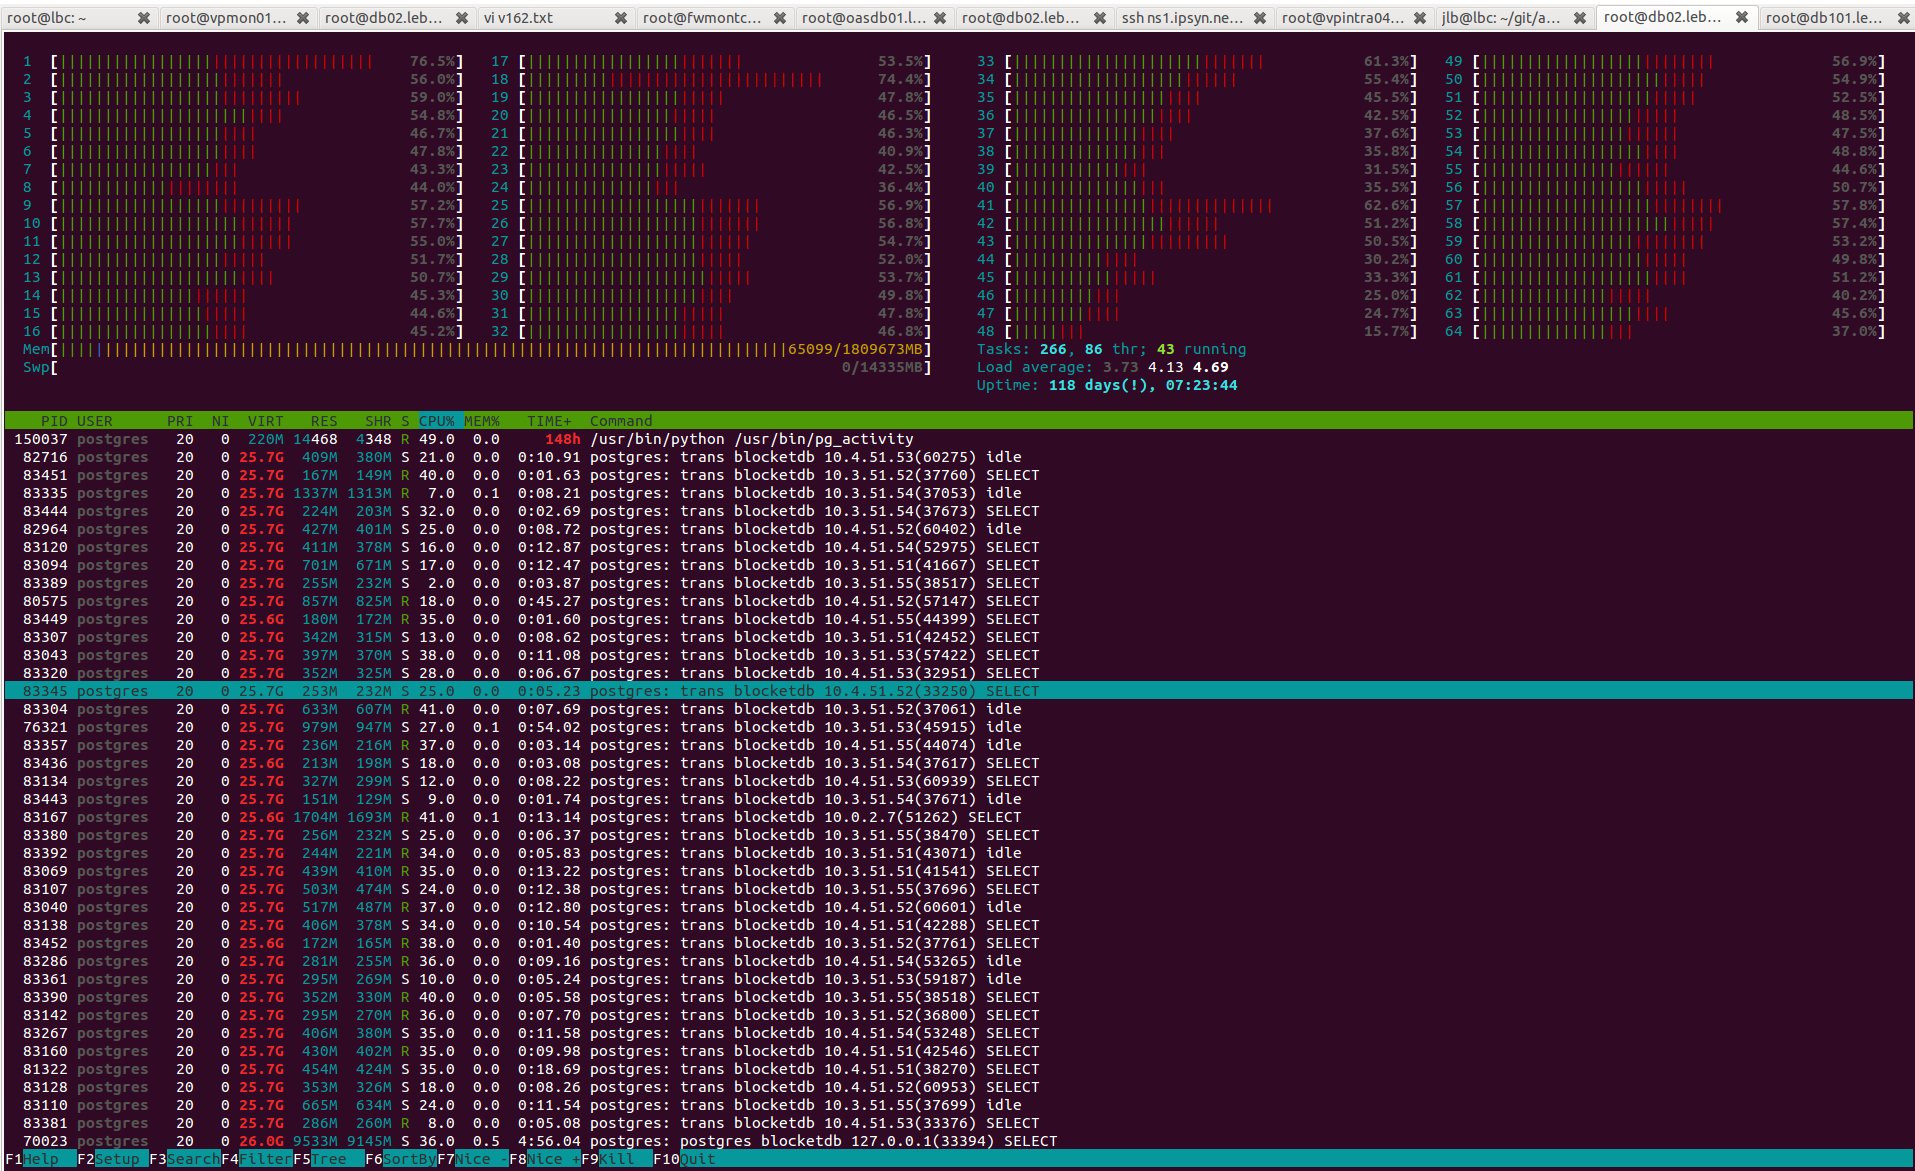
PostgreSQL works beautifully as a DW!



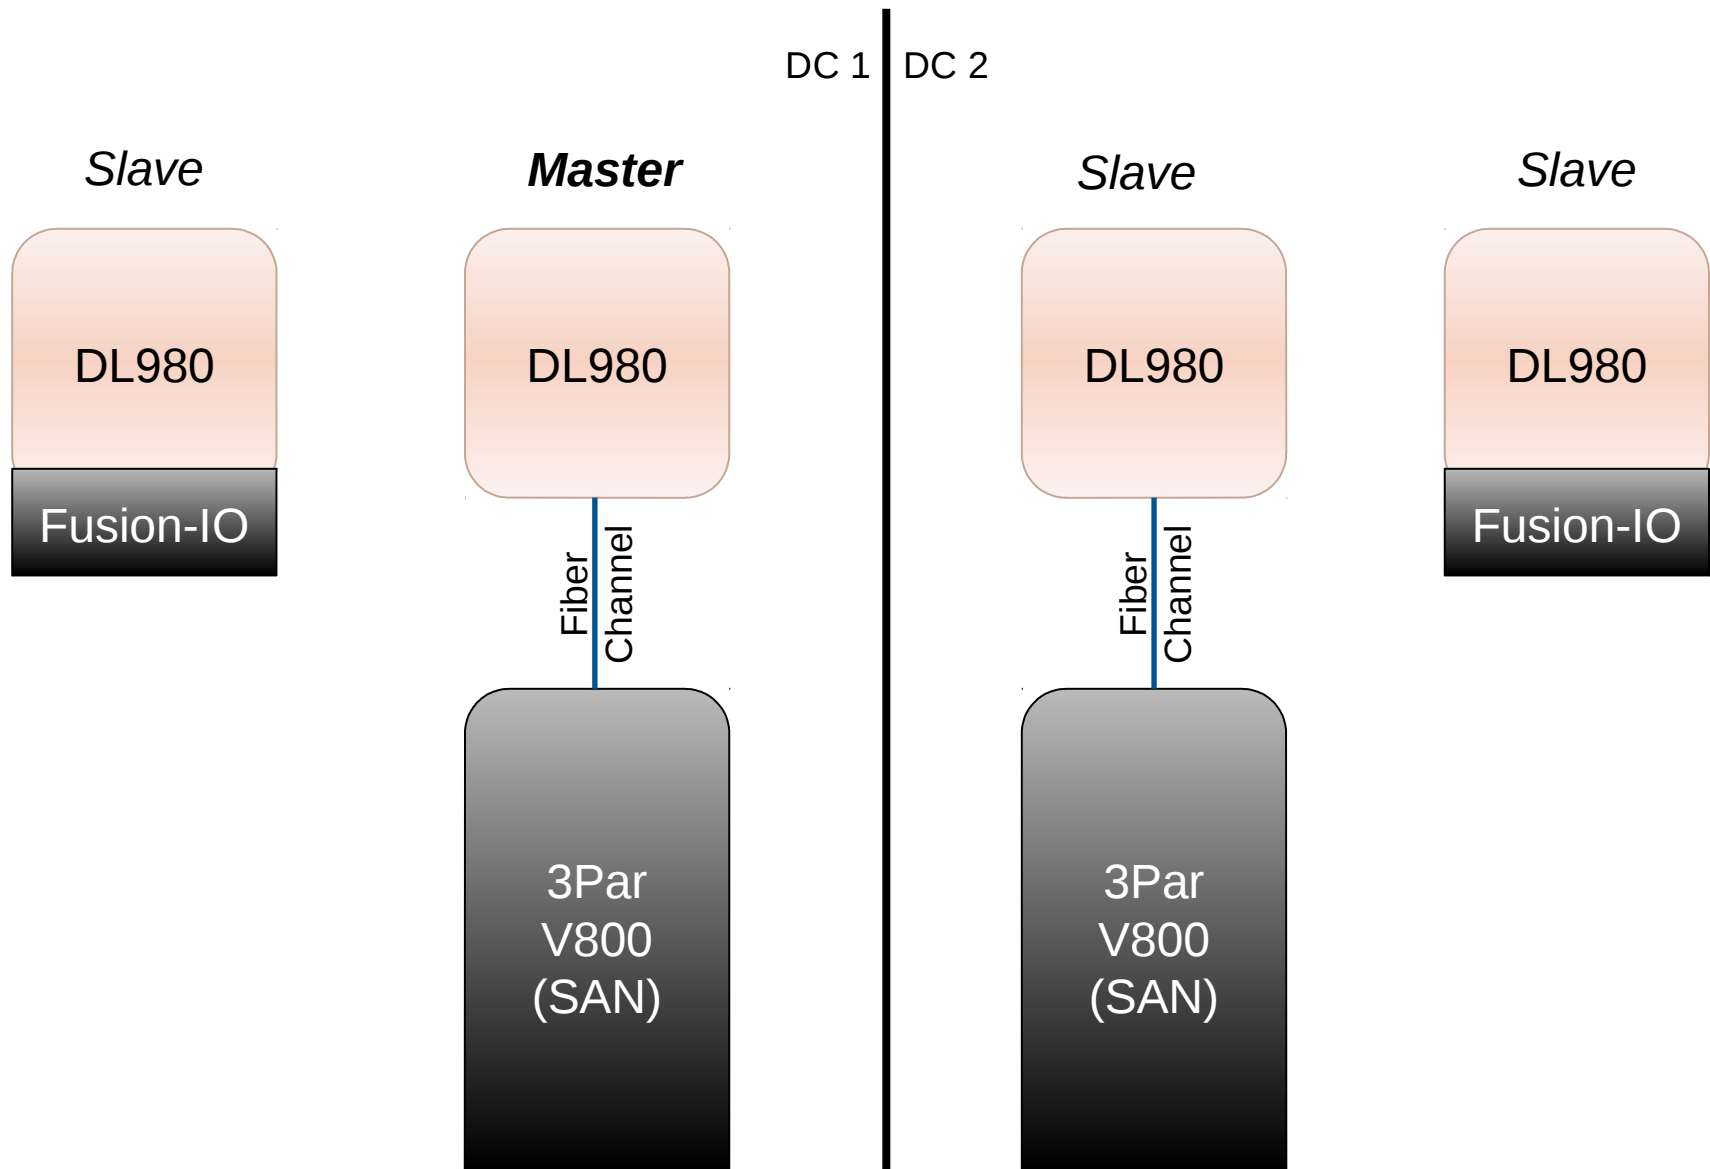
Big Iron: HP DL980, 2TB of RAM, 64-80 cores



HTOP on the Master



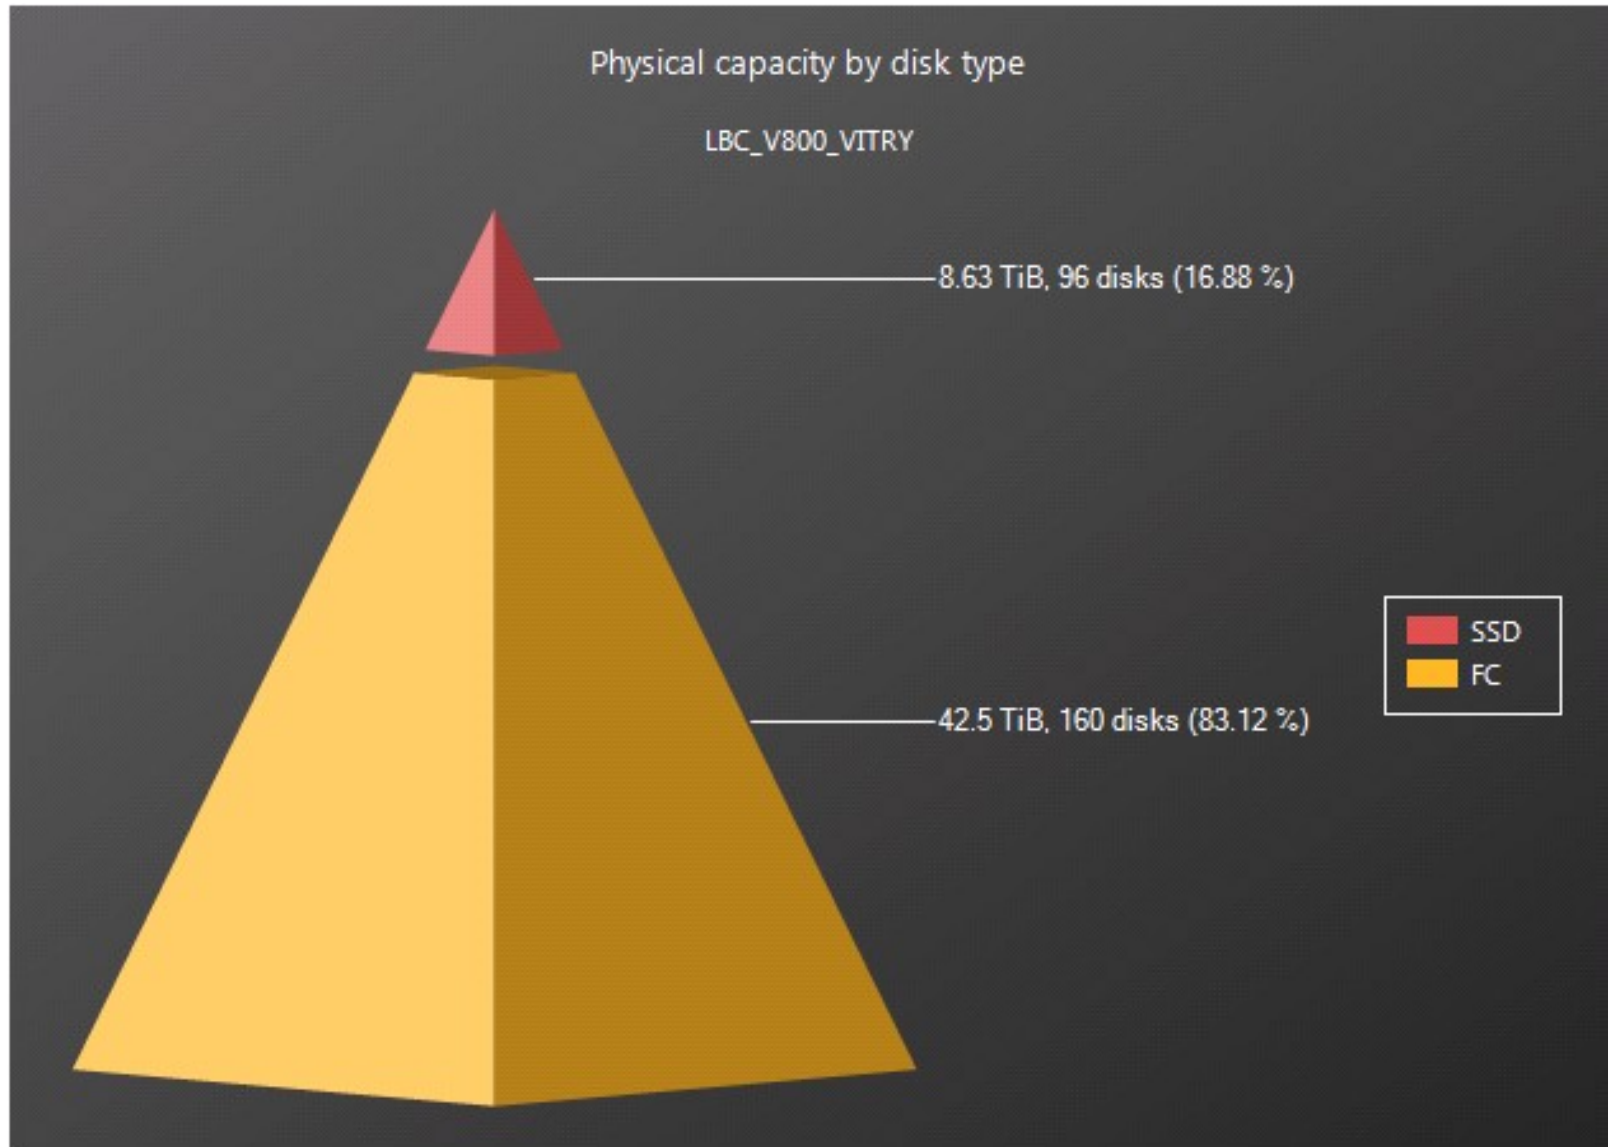
Physical storage for the main PostgreSQL instances



Big Iron: 3Par V800 SAN



Big Iron: thin provisioning with mix of SSD and FC disks



Big Iron: high performance...



Bragging about it ;)

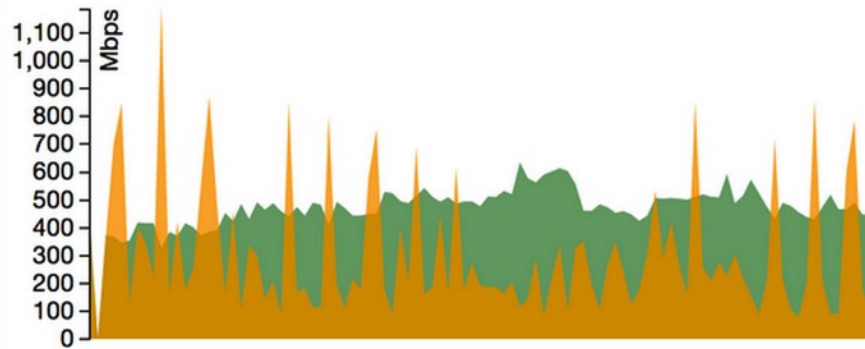


Renaud Bruyeron

@brew_your_own

/proc/diskstats on our largest production database #postgresql #rocks cc @jlb666
pic.twitter.com/BdO2KQAFBK

↩ Reply 🗑 Delete ★ Favorite 📁 Pocket ... More




Bande passante disque

Read **422.6 MBps**


Write **369.8 MBps**

Statistiques /proc/diskstats


Despite the caches and the search engine, we get impressive workloads on the master DB



How did we get there?



We did go through growing pains and near disasters



Our own « worst day of our lives »: March 1st 2013 (1/2)

Master DB is slowing down dramatically

We find that Slony replication is the culprit

We don't know what to do...

Until we find a solution on the net that involves cleaning up slony metadata...

...(you know where this is going)...

We fumble. We notice. The slave is borked.

Rebuilding the Slave with slony brings the Master down. Oh. God...

We take the Master off the stack, and start rebuilding the slave w/ Slony

...5 days later, we are done (!)...

Our own « worst day of our lives »: March 1st 2013 (2/2)

...but we are not out of the woods yet!

Pent up demand is bringing the site down!

We decide to switch to native replication!

...but the network cards are maxed out by the replication data...

...triggering a kernel bug...

...(Murphy, could you please step out of the room?)...

We implement network card bonding, and start moving support tables off the main instance

...and we are done!



SCHIBSTED
CLASSIFIED MEDIA

What's Next?


Vertical Scalability has limits!

We are already running on the biggest HW money can buy

Past certain volume levels, execution plans can change radically

Huge instances are difficult to maintain & backup safely

Rebuilding the slave in March 2013 took a full 5 days...



Although we are maxing out the HW available to us, especially on the writes/s,
We are committed to PostgreSQL at the heart of our platform

Key ideas to take our platform to the next level

Sharding



Spread the reads & writes horizontally

Unbundling of
the schema



Move parts of the schema (that can be decoupled) to other instances. Spread the workload

PGQ



Reduce application-level transaction interleaving by moving parts of the transactions to asynchronous workers

Questions?

