



PGConf.Brasil 2018



Zero Downtime PostgreSQL Upgrades

Erivelton Vichroski



Informação pública

Digitro
40 anos

 eriveltonvichroski
erivelton.vichroski@digitro.com.br

select * from eu;

Graduado em Sistemas de Informação pela UFSC;

Certificações em Oracle, DB2 e MongoDB;

PostgreSQL, Oracle, DB2, SQLServer, Informix, Linux, Windows Server, Migrações e Integrações de dos dados, Java, Python, etc..

Desafio atual: Big Data





CRENCIADA COMO EMPRESA
ESTRATÉGICA DE DEFESA

TECNOLOGIA BRASILEIRA

Digitro

40
anos

CRIANDO TECNOLOGIAS PARA
UM MUNDO MELHOR!

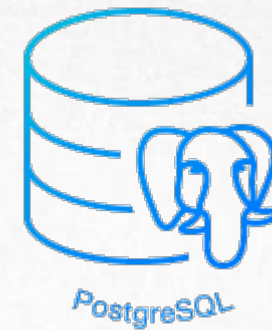


Digitro
40 anos



PGConf.Brasil 2018

Agenda

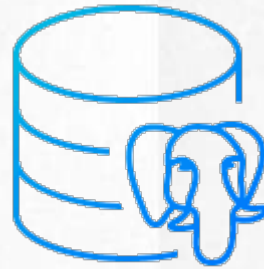


▣ Motivação;

▣ Tecnologias Envolvidas;

▣ Proposta de Solução;

▣ Considerações Finais.



PostgreSQL

Motivação



O Problema

- ▣ Prover aplicações SaaS/Multi tenant;
- ▣ Requisitos de *downtime* (tempo de indisponibilidade) **quase zero**;

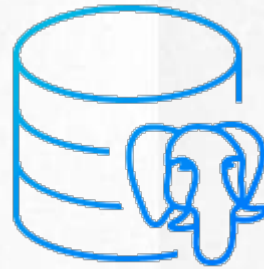
- ▣ Garantir SLA.



% DISPONIBILIDADE	HORAS ANO 24x7	TEMPO MÁXIMO DE INDISPONIBILIDADE	HORAS TOTAIS ANO 8x5	TEMPO MÁXIMO DE INDISPONIBILIDADE
90%	8760horas	876horas 36,5 dias	2086horas	208.6horas 26 dias
95%	8760horas	438horas 18,3 dias	2086horas	104horas 13 dias
99%	8760horas	87horas 3,6 dias	2086horas	20.86horas 2,6 dias
99,9%	8760horas	8,76horas	2086horas	2horas
99,982%	8760horas	52,56minutos	2086horas	12,5minutos
99,999%	8760horas	5,3minutos	2086horas	1,25minutos

O Problema

- ▣ Necessidade de atualizações frequentes do PostgreSQL:
 - > segurança;
 - < risco de corrupção de dados;
 - > performance.
- ▣ Algumas de nossas APIs não toleram 500 *Internal Server Error* **post** *pj/dados* **HTTP/1.1**

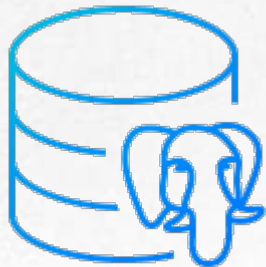


PostgreSQL

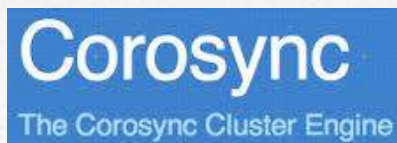
Proposta de Solução



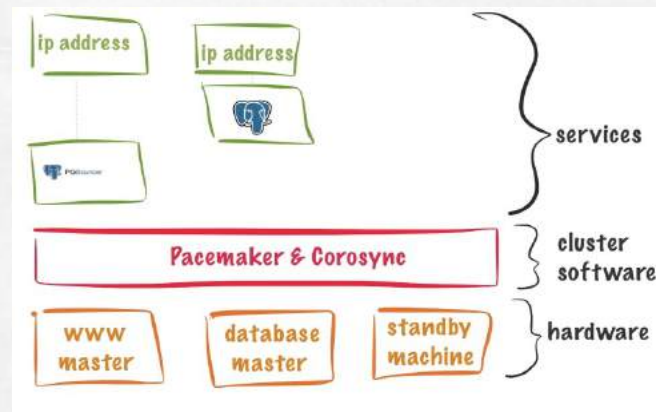
Tecnologias Envolvidas



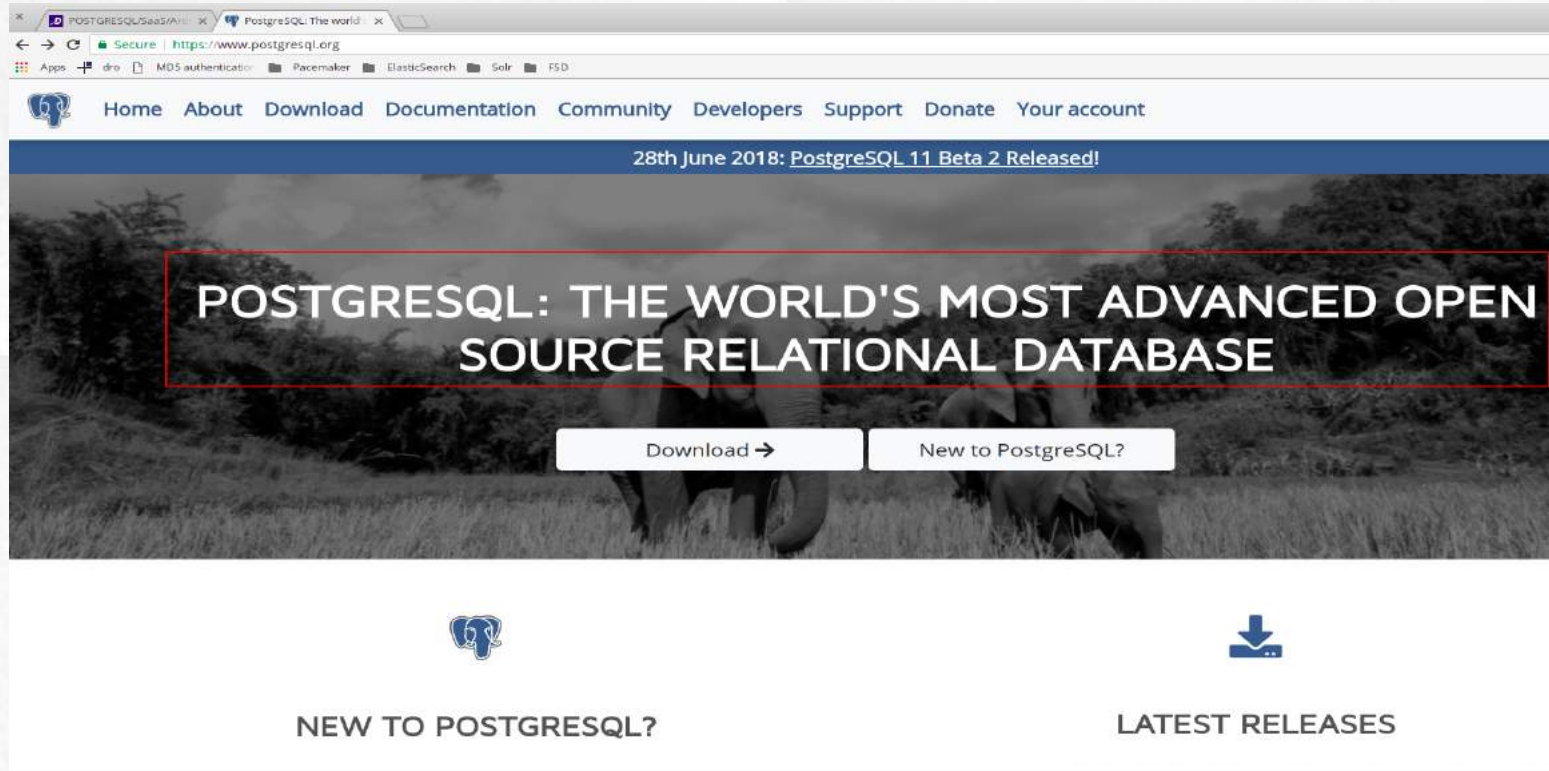
PostgreSQL



PGBouncer



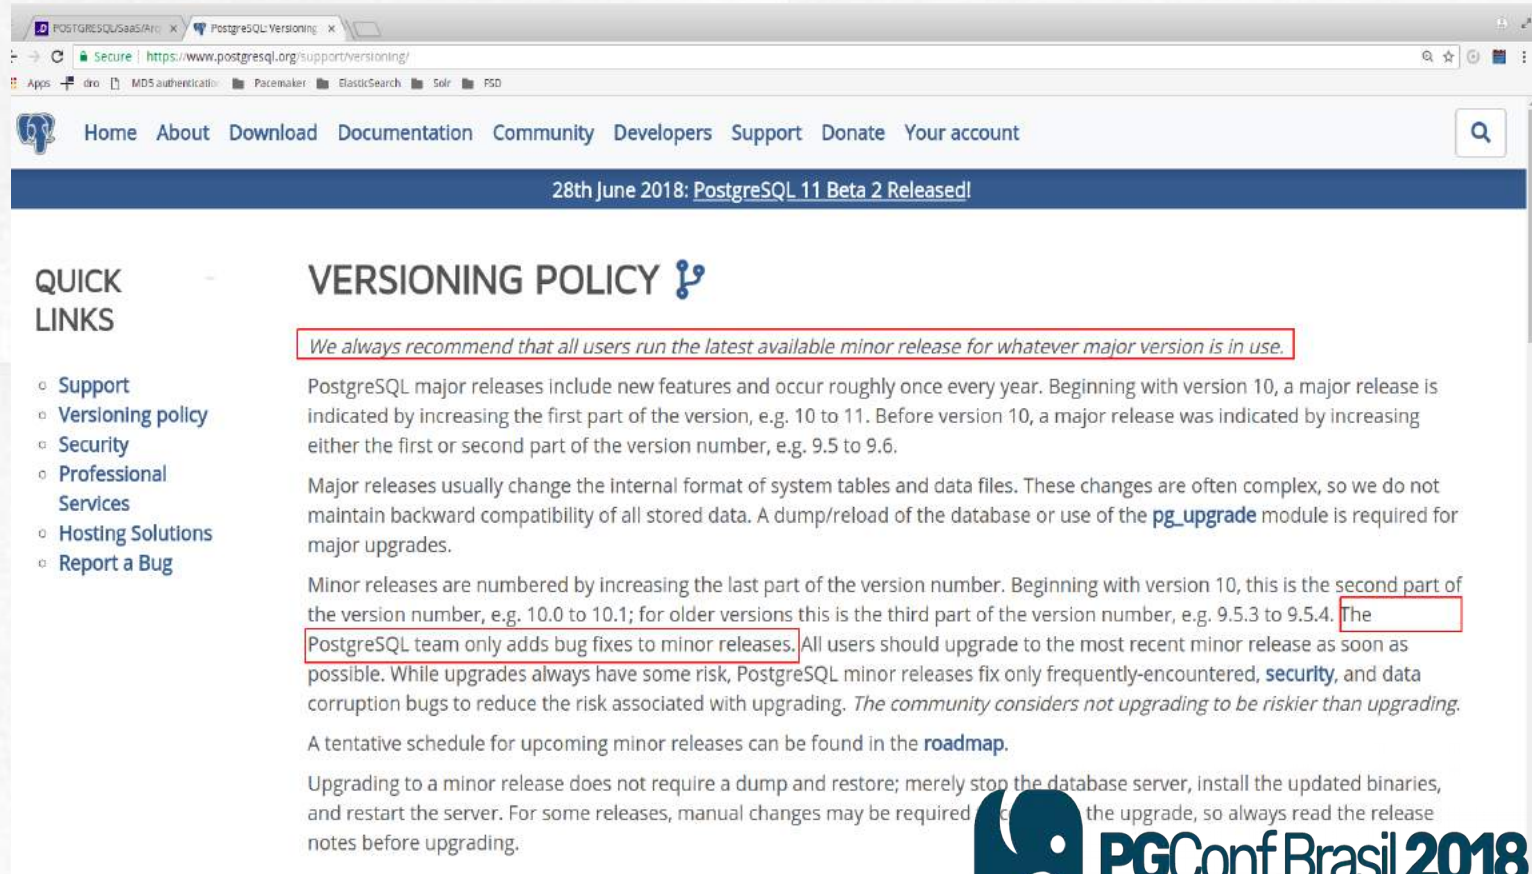
PostgreSQL



The image shows a screenshot of the PostgreSQL website homepage. At the top, there is a navigation menu with links for Home, About, Download, Documentation, Community, Developers, Support, Donate, and Your account. Below the navigation is a blue banner with the text "28th June 2018: PostgreSQL 11 Beta 2 Released!". The main content area features a large image of elephants in a savanna, with a red-bordered box containing the text "POSTGRESQL: THE WORLD'S MOST ADVANCED OPEN SOURCE RELATIONAL DATABASE". Below this text are two buttons: "Download →" and "New to PostgreSQL?". At the bottom of the page, there are two columns: "NEW TO POSTGRESQL?" with a PostgreSQL logo icon, and "LATEST RELEASES" with a download icon.



Upgrade do PostgreSQL



The screenshot shows the PostgreSQL website's versioning policy page. The browser address bar displays the URL <https://www.postgresql.org/support/versioning/>. The navigation menu includes links for Home, About, Download, Documentation, Community, Developers, Support, Donate, and Your account. A blue banner at the top of the page reads "28th June 2018: PostgreSQL 11 Beta 2 Released!".

QUICK LINKS

- Support
- Versioning policy
- Security
- Professional Services
- Hosting Solutions
- Report a Bug

VERSIONING POLICY

We always recommend that all users run the latest available minor release for whatever major version is in use.

PostgreSQL major releases include new features and occur roughly once every year. Beginning with version 10, a major release is indicated by increasing the first part of the version, e.g. 10 to 11. Before version 10, a major release was indicated by increasing either the first or second part of the version number, e.g. 9.5 to 9.6.

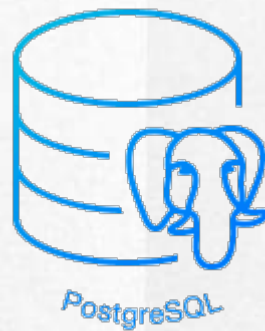
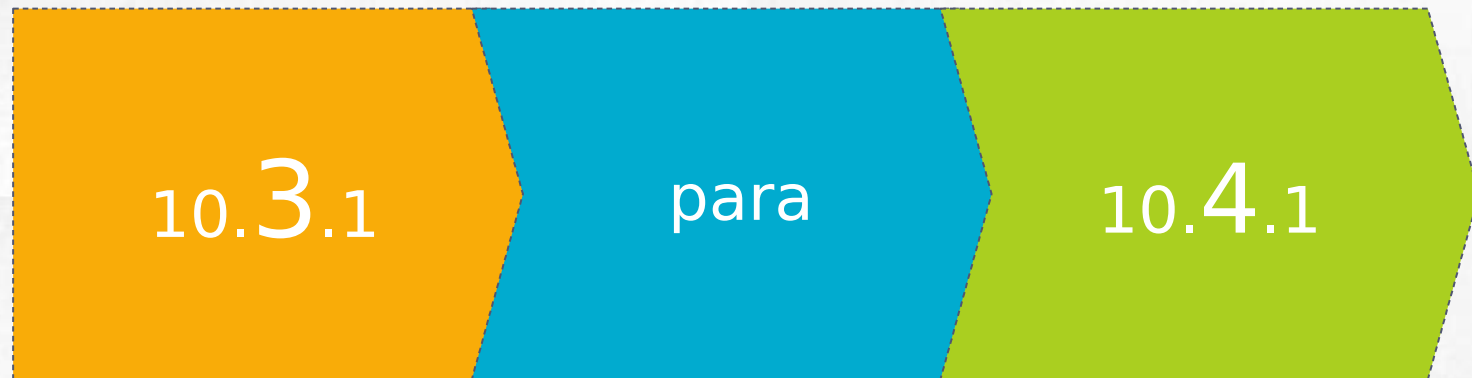
Major releases usually change the internal format of system tables and data files. These changes are often complex, so we do not maintain backward compatibility of all stored data. A dump/reload of the database or use of the `pg_upgrade` module is required for major upgrades.

Minor releases are numbered by increasing the last part of the version number. Beginning with version 10, this is the second part of the version number, e.g. 10.0 to 10.1; for older versions this is the third part of the version number, e.g. 9.5.3 to 9.5.4. The PostgreSQL team only adds bug fixes to minor releases. All users should upgrade to the most recent minor release as soon as possible. While upgrades always have some risk, PostgreSQL minor releases fix only frequently-encountered, **security**, and data corruption bugs to reduce the risk associated with upgrading. *The community considers not upgrading to be riskier than upgrading.*

A tentative schedule for upcoming minor releases can be found in the [roadmap](#).

Upgrading to a minor release does not require a dump and restore; merely stop the database server, install the updated binaries, and restart the server. For some releases, manual changes may be required [before](#) the upgrade, so always read the release notes before upgrading.

Upcoming minor releases



Implicações na Disponibilidade

▣ Exemplo: subir da 10.3.1 para versão 10.4.1;

▣ Dump/Restore geralmente **não é requerido**;

▣ Indisponibilidade:
(atualizar binários + reiniciar)

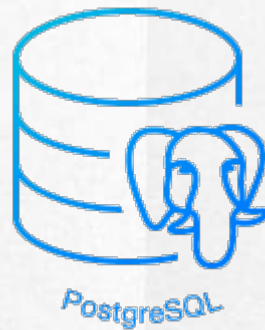
<https://www.postgresql.org/docs/9.6/static/release-9-6-4.html>

Next major releases

9.6.9

para

10.4.1



Implicações na Disponibilidade

▣ Exemplo: subir de 9.6.9 para versão 10.4.1;

▣ Dump/Restore ou pg_upgrade **é requerido;**

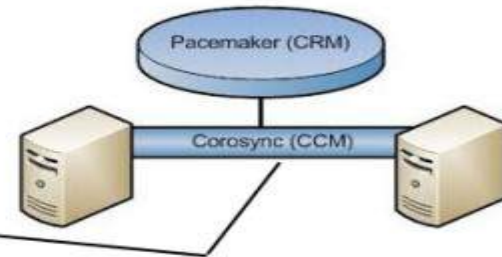
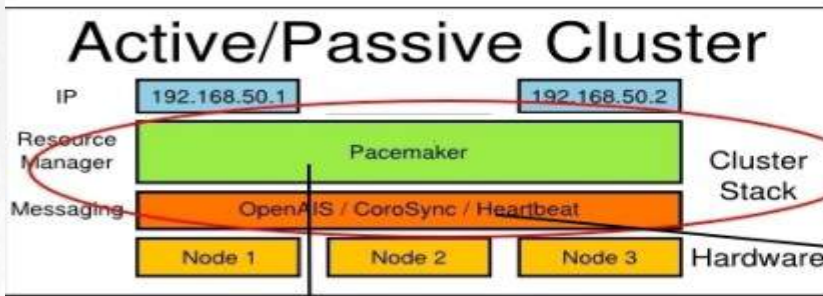
▣ Tempo > de indisponibilidade
(Upgrade de binários + catálogo + bancos).

Pacemaker



- ▣ É um CRM (Cluster Resource Manager);
- ▣ Fornece alta disponibilidade para serviços;
- ▣ Serviço de Heartbeat /VIP (com corosync).

Pacemaker

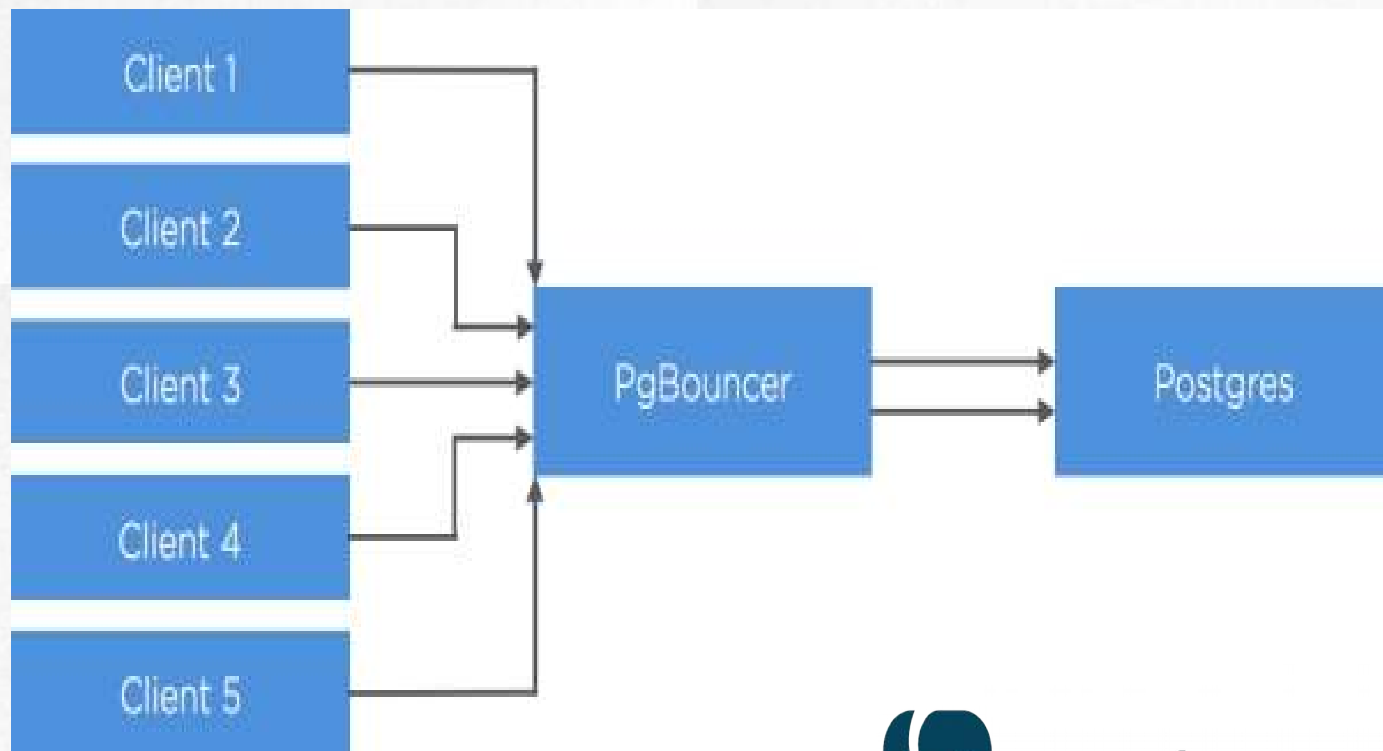


PGBouncer

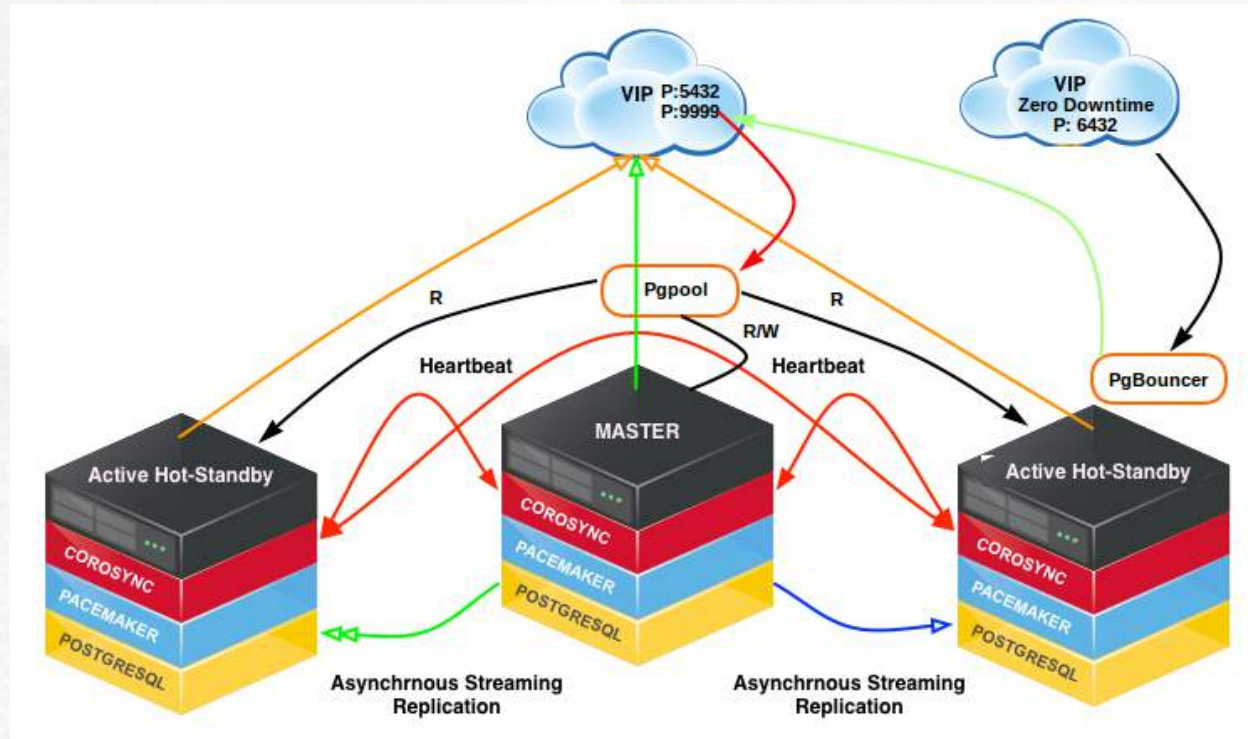


- ▣ Agrupador de conexões para o PostgreSQL (Proxy);
- ▣ Apps se conectam ao PGBouncer “pensando” que é o servidor PostgreSQL;
- ▣ Pool de Conexão (reuso de conexão);
- ▣ *Pgbouncer does not fork a new process (it is event-based).*

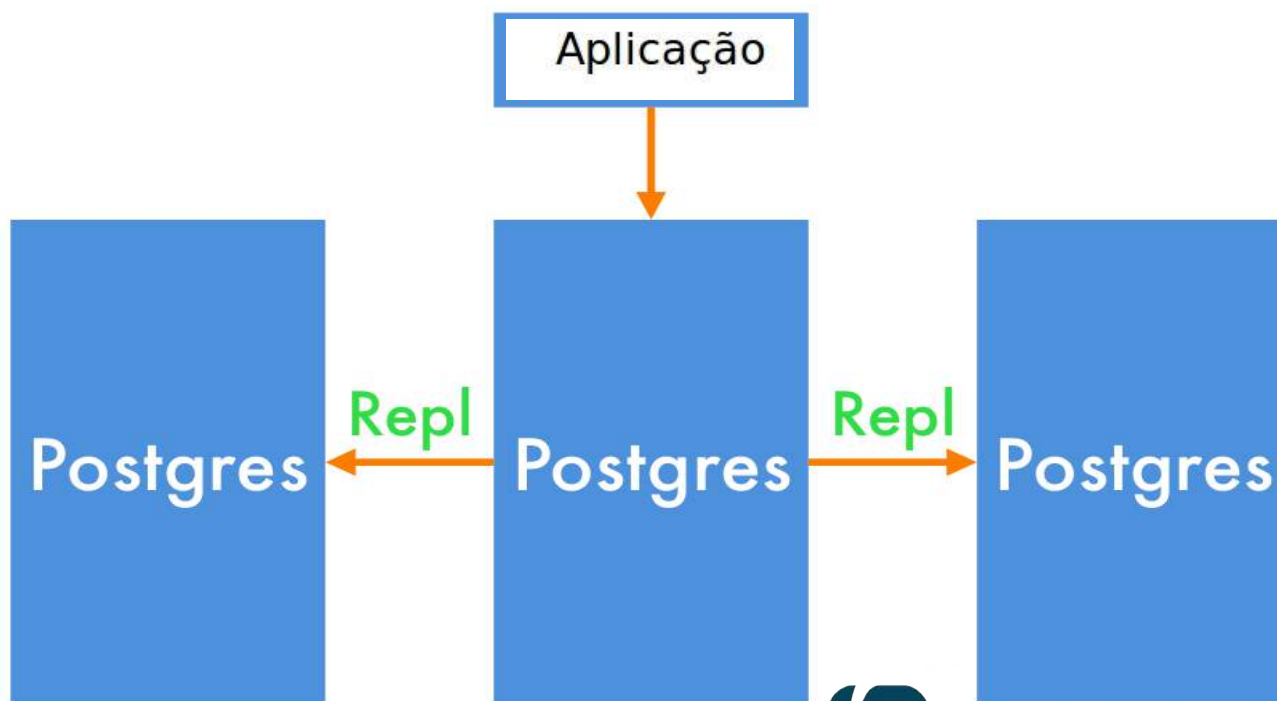
PGBouncer



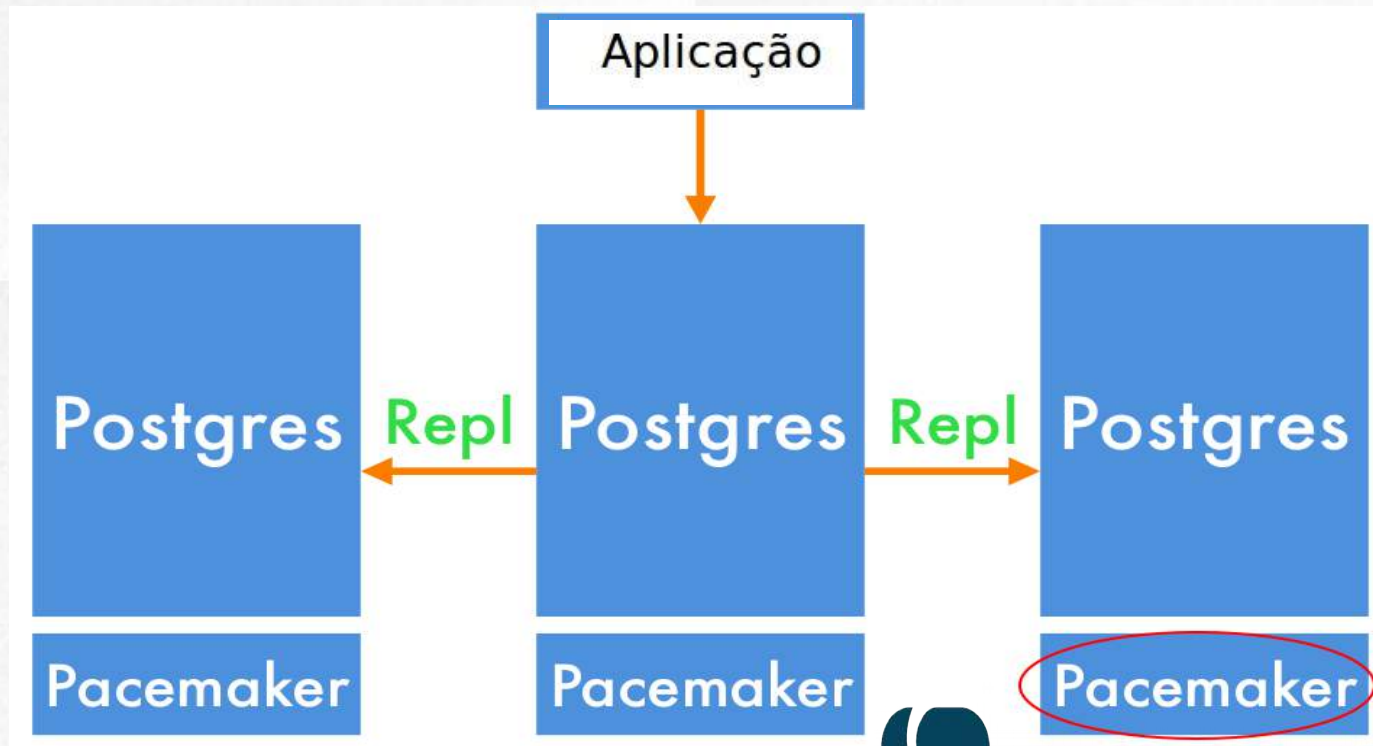
Proposta de Solução



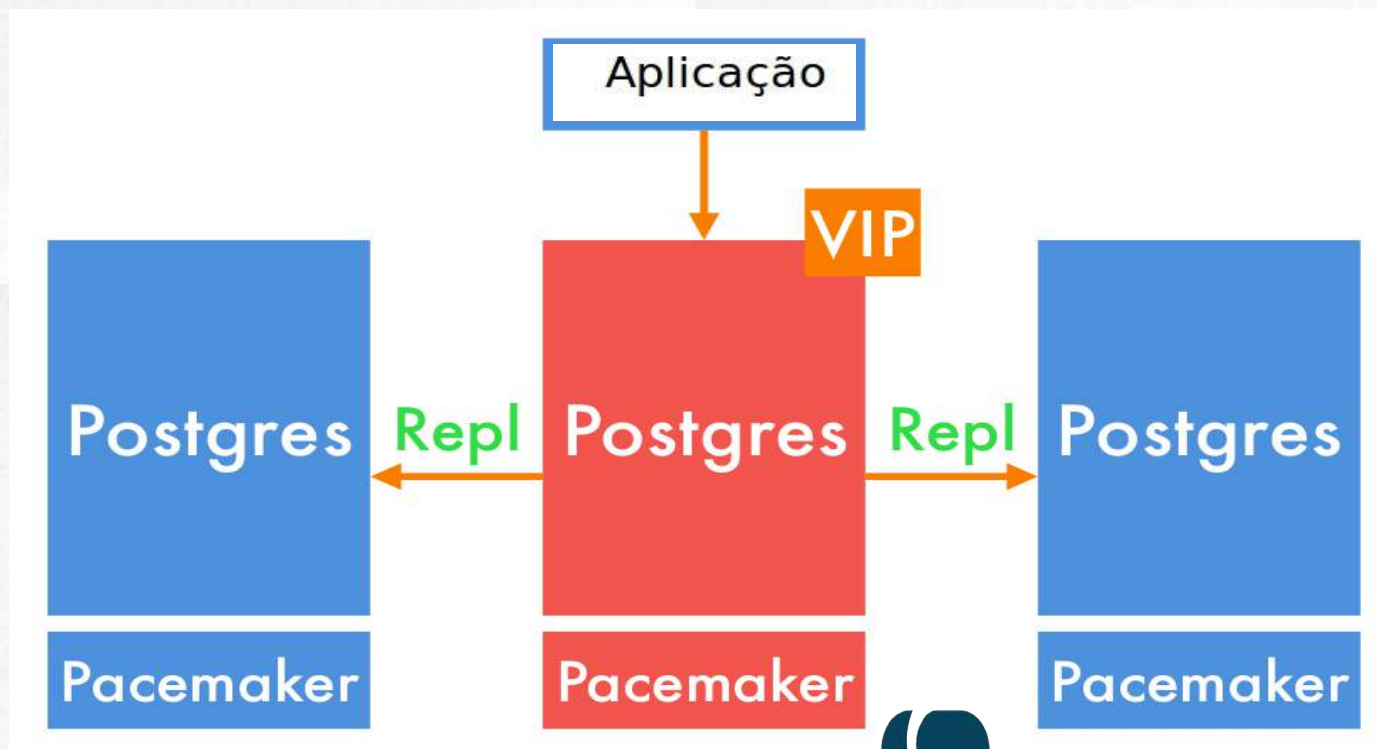
Switchover sem Zero Downtime



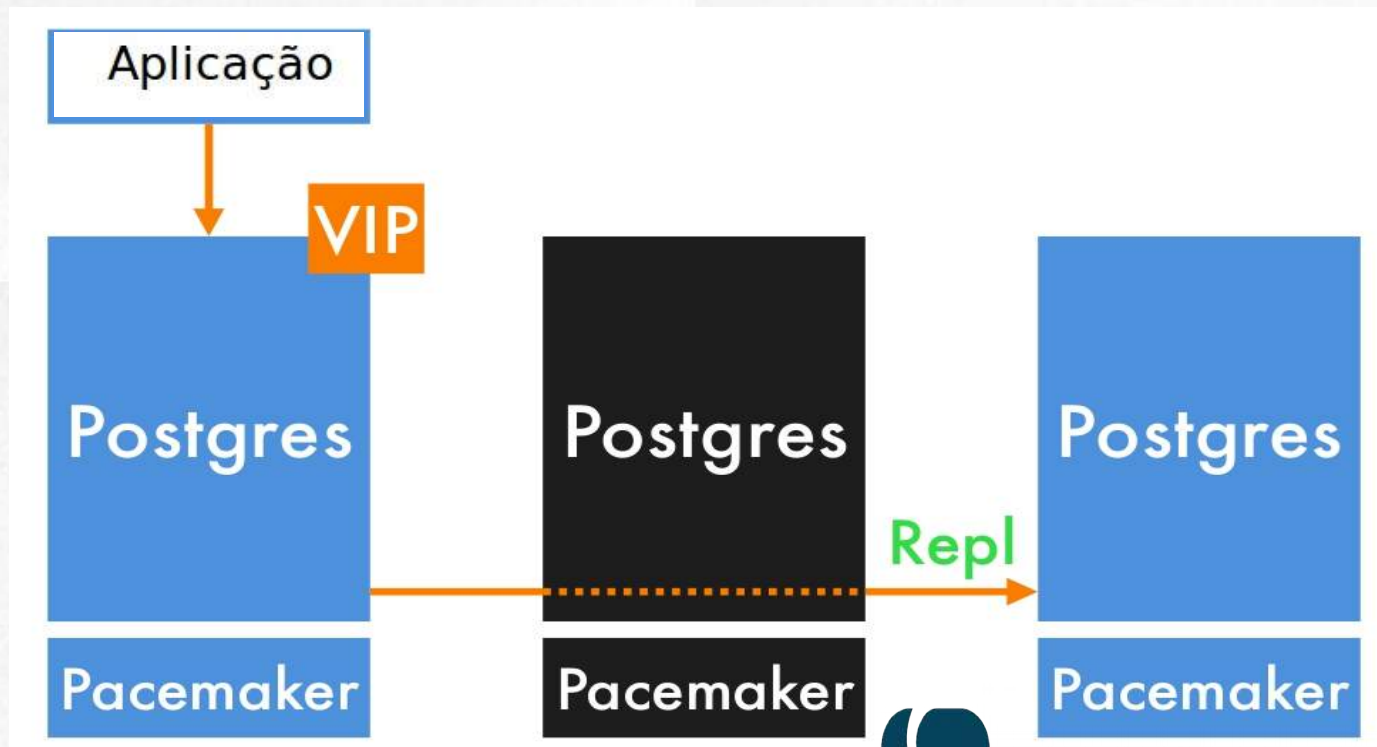
Switchover sem Zero Downtime



Switchover sem Zero Downtime



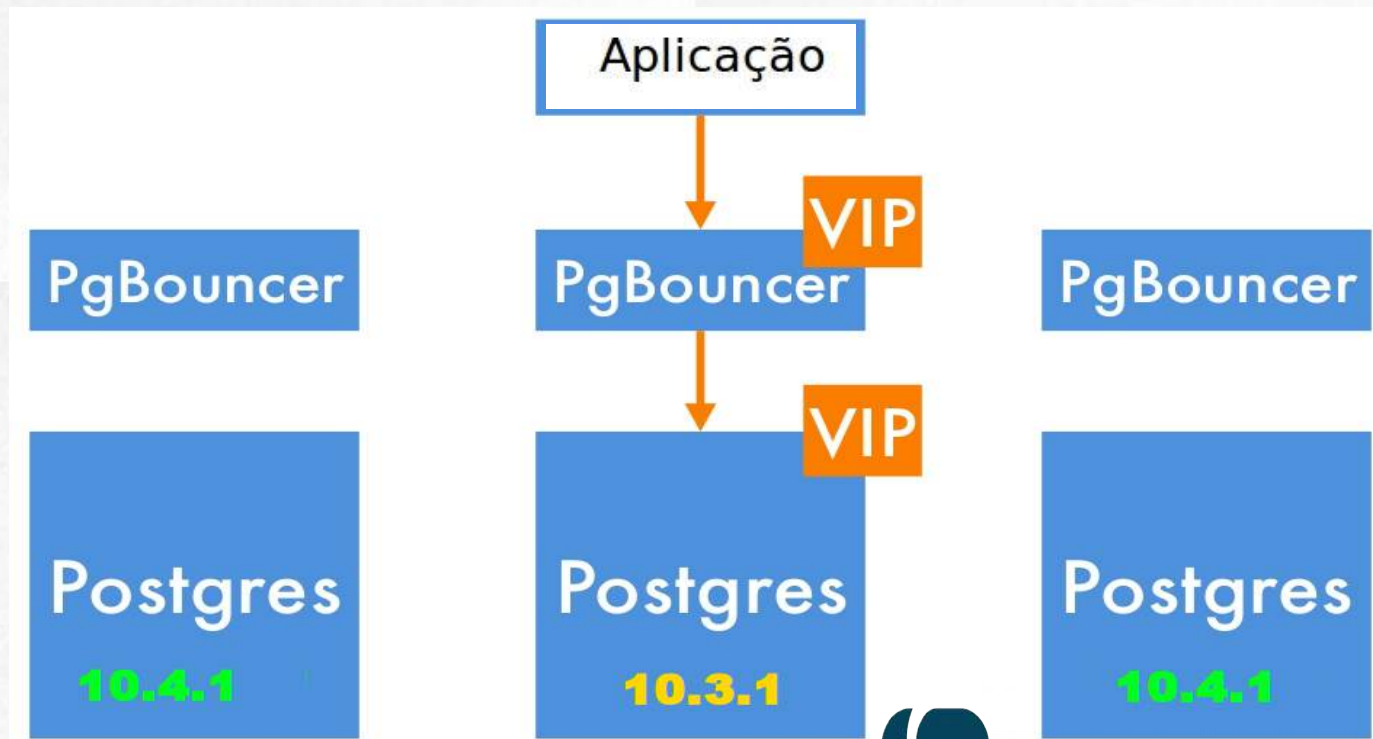
Switchover sem Zero Downtime



Switchover com Zero Downtime



Switchover com Zero Downtime



*Switchover **com** Zero Downtime*

PAUSE;



PGBouncer



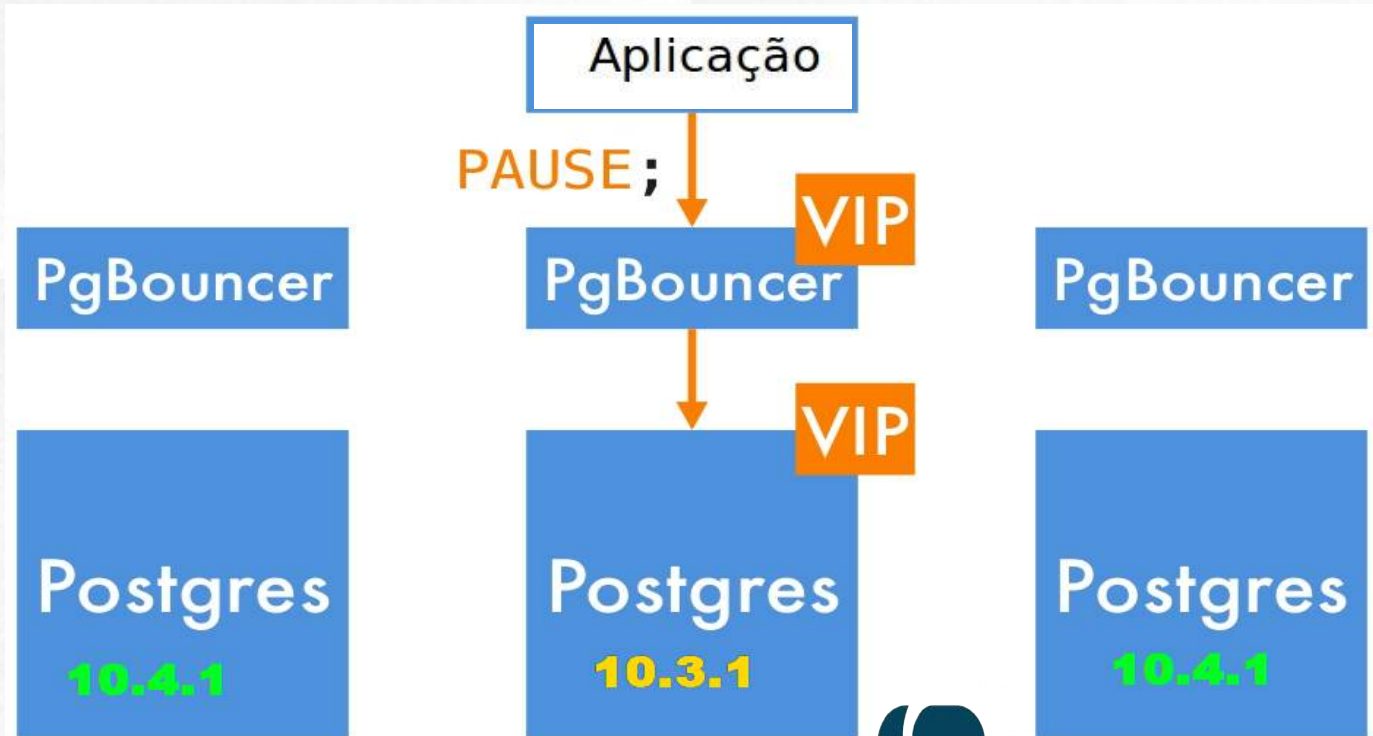
PGConf.Brasil 2018

PGBouncer

PAUSE;

“Segura” todas as instruções SQL de entrada destinada ao PGBouncer colocando em fila.

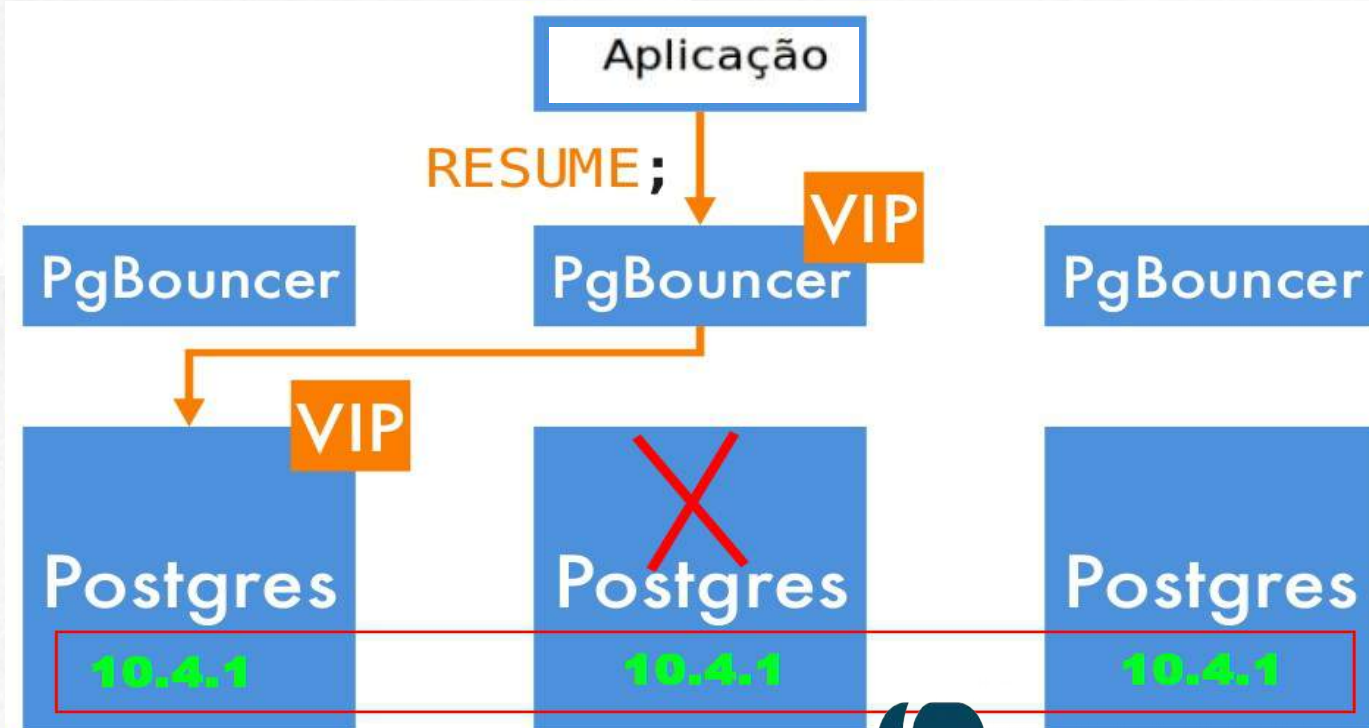
Switchover com Zero Downtime



Switchover com Zero Downtime



Switchover com Zero Downtime



Considerações sobre a Arquitetura Zero Downtime

- ▣ *Minor version* (PostgreSQL);
- ▣ *Long-running transactions*;
- ▣ PAUSE - PGBouncer ~ 10 secs.

Prova de Conceito

▣ Virtualbox;

▣ Vagrant;

▣ <https://github.com/gocardless/our-postgresql-setup>

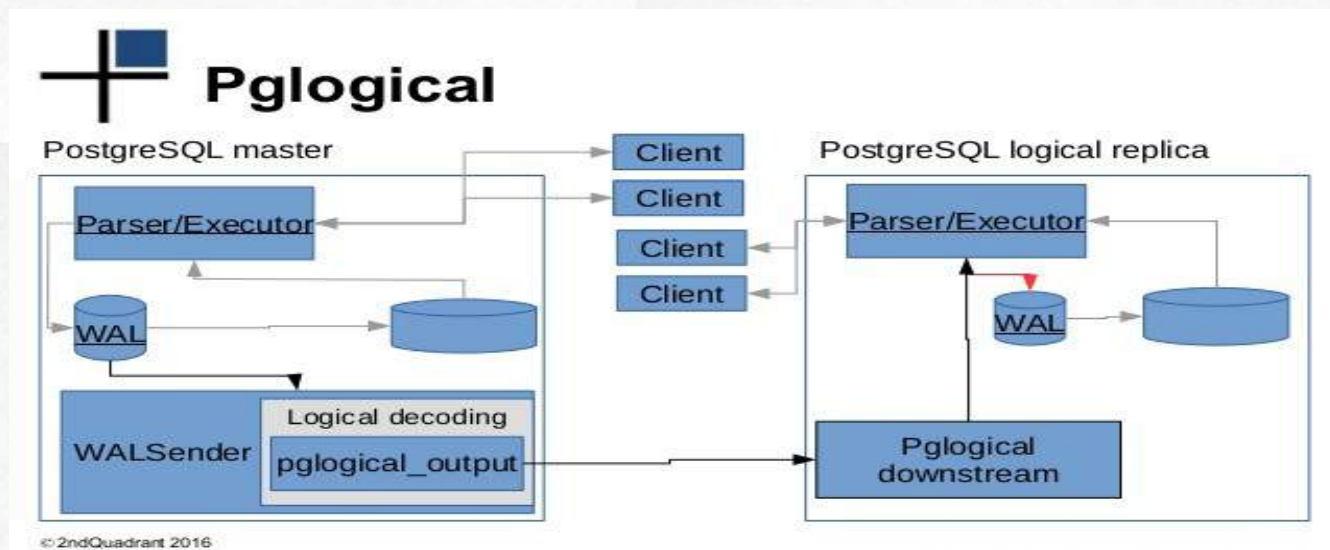
Considerações Finais

Desafios e trabalhos futuros:

- ▣ Migrar com *Near-Zero Downtime* o PostgreSQL de 10.x para a 11.x;
- ▣ Ambiente exige I/O elevado de R/W com requisito de baixa latência.
Migração não pode impactar o ambiente.

Considerações Finais

- ▣ Upgrade major version com < indisponibilidade:



Considerações Finais

- ▣ *Upgrade major version com < indisponibilidade:*

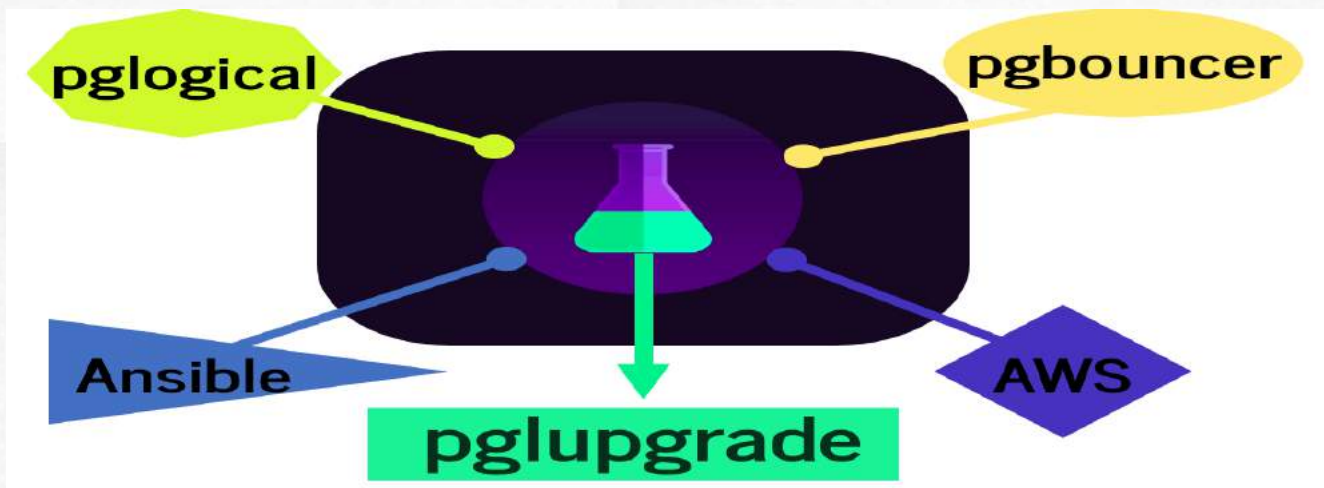


Imagem extraída da apresentação
Near-Zero Downtime Automated Upgrades of PostgreSQL
Clusters in Cloud . Gülçin Yıldırım – 2ndQuadrant
<http://slides.com/apatheticmagpie/automated-upgrades-of-postgres-2#/10>

Considerações Finais

▣ Upgrade major version com < indisponibilidade:

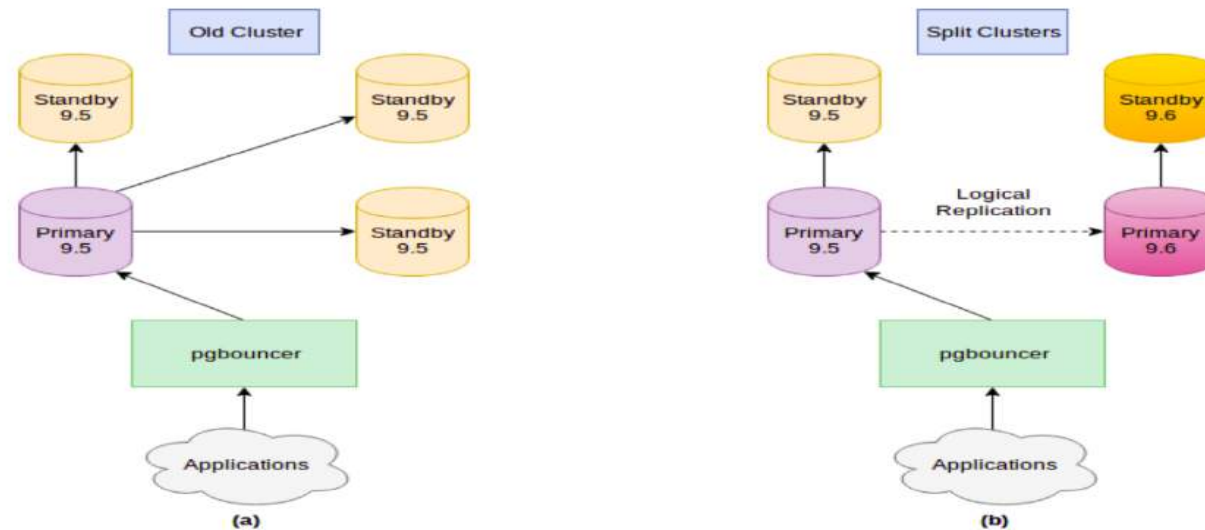


Imagem extraída da apresentação
Near-Zero Downtime Automated Upgrades of PostgreSQL
Clusters in Cloud . Gülçin Yıldırım – 2ndQuadrant
<http://slides.com/apatheticmagpie/automated-upgrades-of-postgres-2#/14>

Considerações Finais

Código:

<https://github.com/eriveltonvichroski/postgresql-zero-down-upgrades>

Obrigado!
Perguntas
??

