

AN ILLUSTRATED HISTORY OF

POSTGRES(QL)

PETER VAN HARDENBERG

- ▶ Started Heroku Postgres, which runs >1M Postgres customer databases.
- ▶ Alright at writing SQL.
- ▶ Started with PostgreSQL in 2001 for arctic oceanography data.
- ▶ I've read the whole Postgres manual.
- ▶ Born in Canada, live in San Francisco.
- ▶ [@pvh](#) / pvh@pvh.ca





INTRODUCTION

MEET
POSTGRES

MEET POSTGRES

**BORN IN CALIFORNIA
IN 1985**

MEET POSTGRES

**BUILT BY A GLOBAL
COMMUNITY SINCE 1996**

MEET POSTGRES

**TRUSTED BY COMPANIES
AROUND THE WORLD**

CHAPTER ONE

STONEBRAKER'S POSTGRES



BERKELEY, CALIFORNIA, 1976

**MICHAEL STONEBRAKER
IS WORRIED**

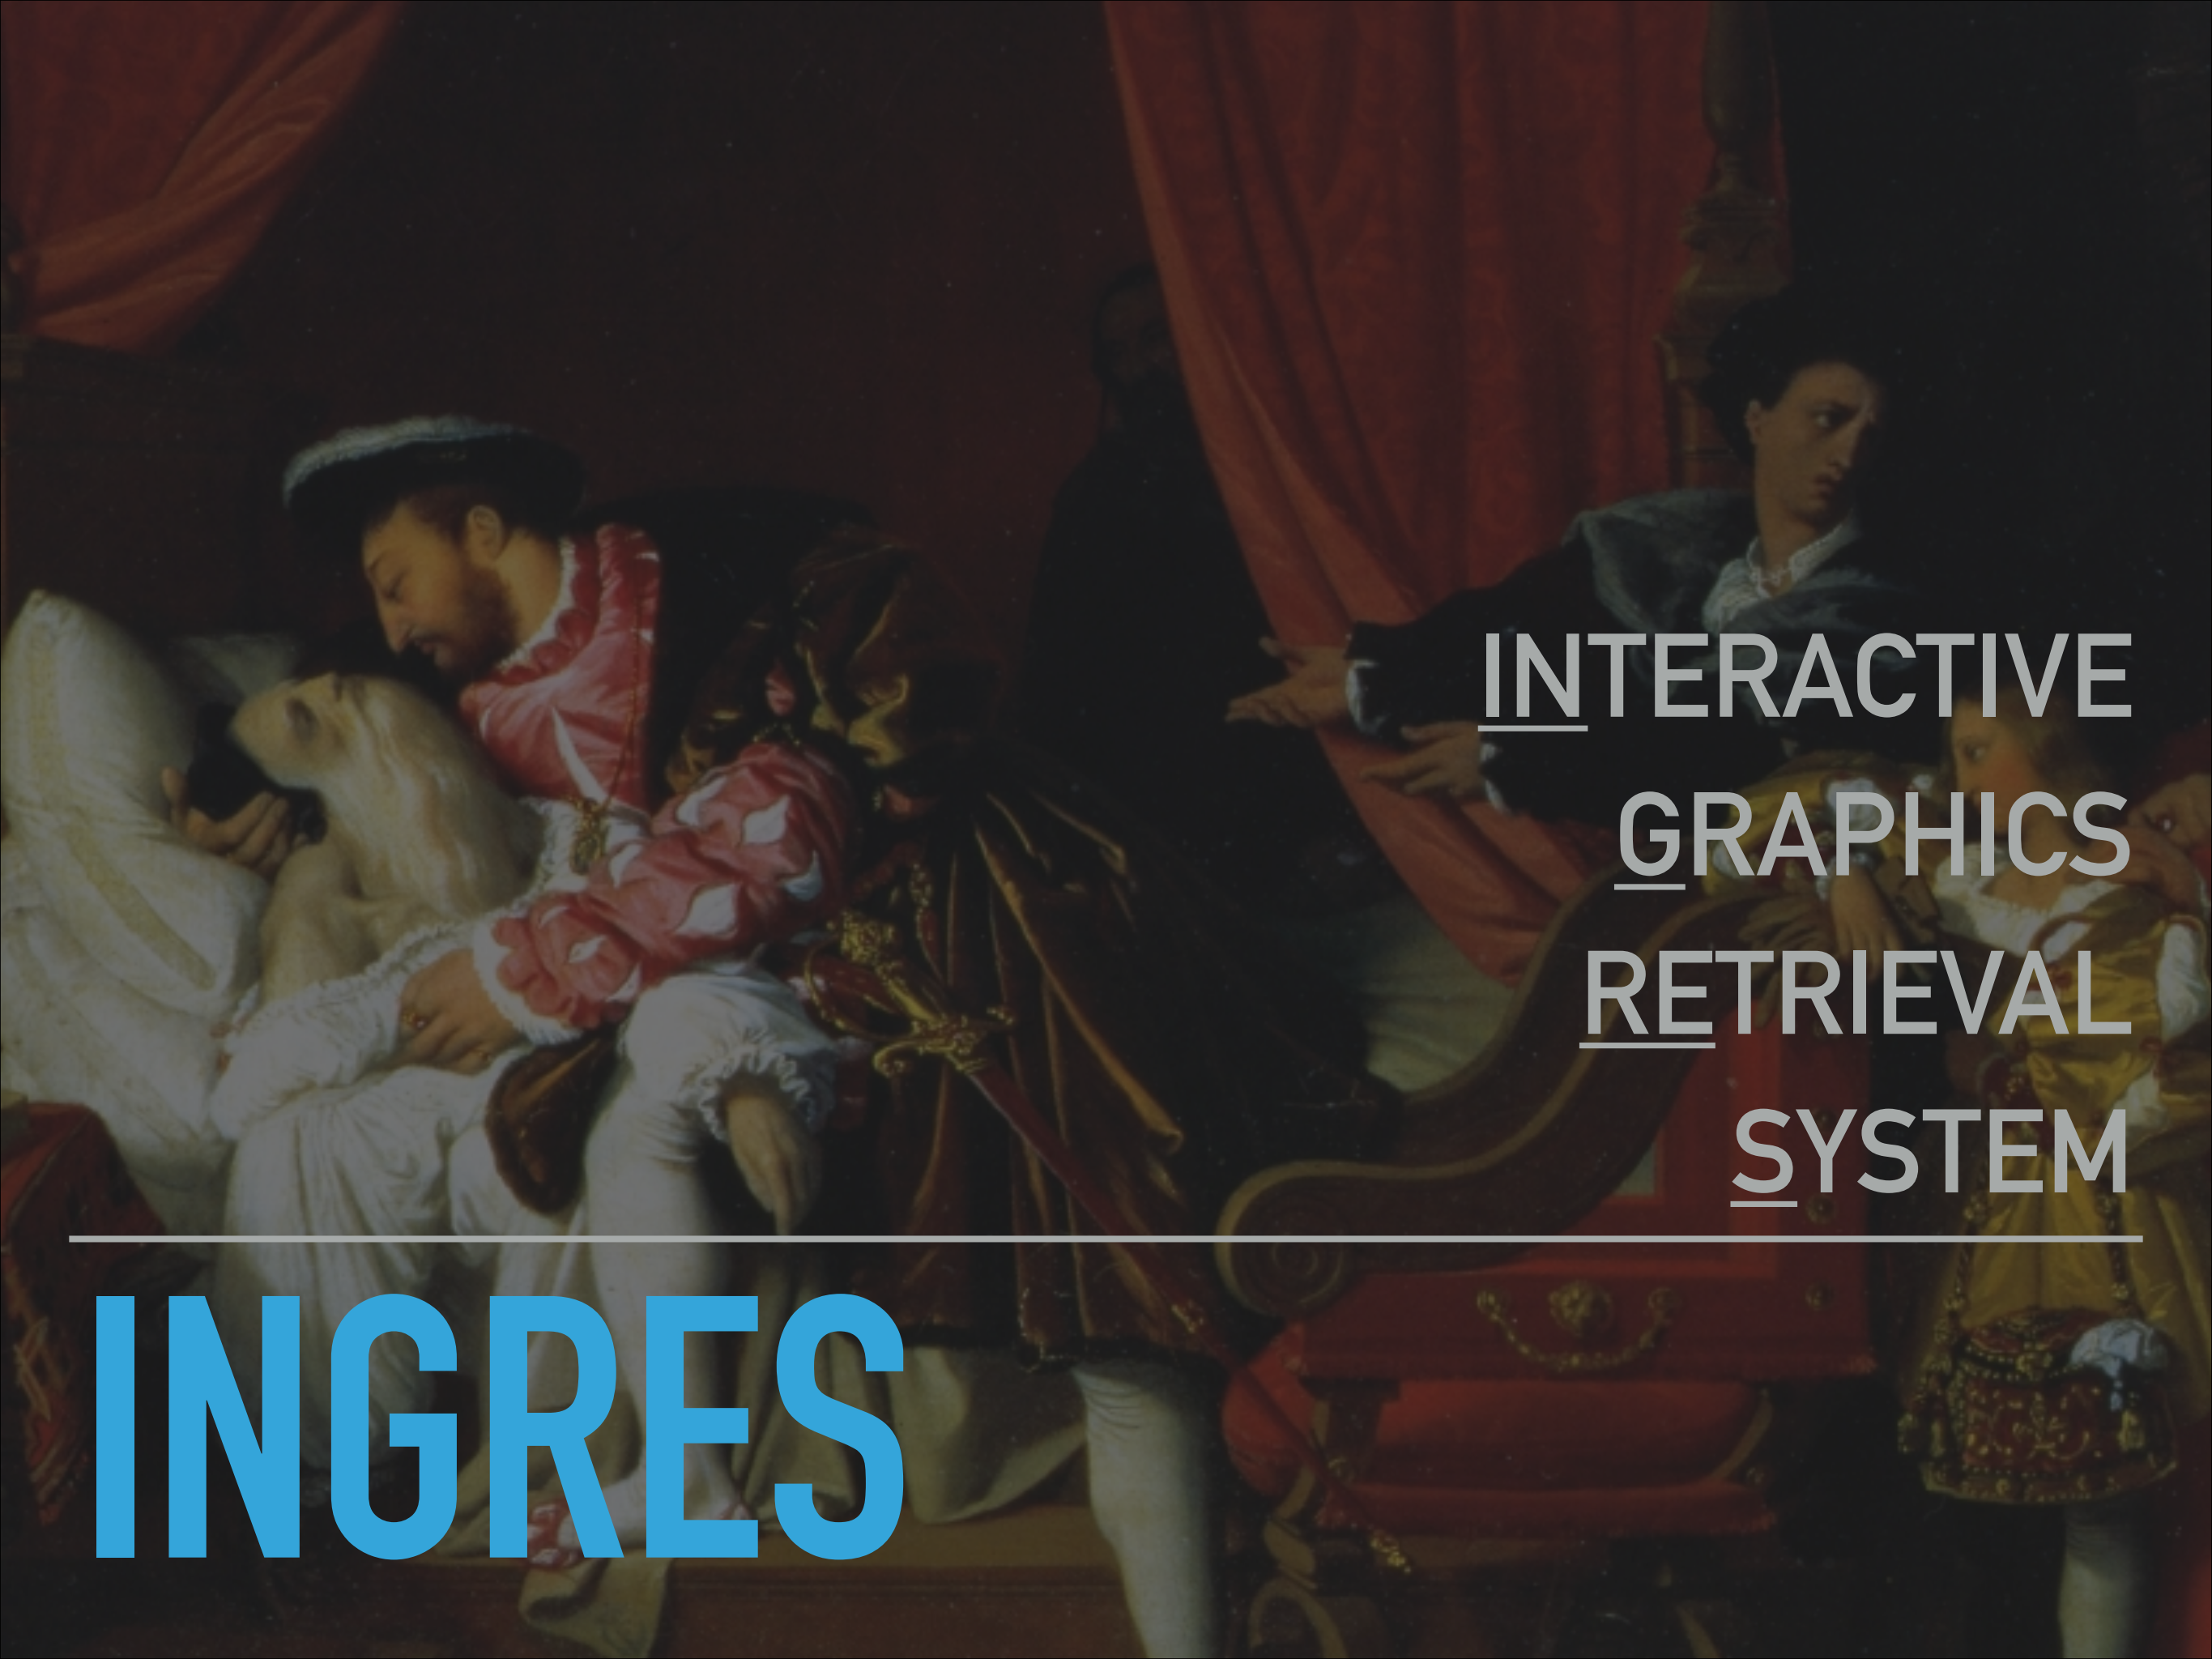
**ONCE YOU GET IN, YOU'VE GOT
FIVE YEARS TO GET FAMOUS,
THAT'S IT.**

Michael Stonebraker



EDGAR F. CODD

**RELATIONAL
ALGEBRA**



INTERACTIVE
GRAPHICS
RETRIEVAL
SYSTEM

INGRES

range of E is EMPLOYEE

retrieve into W

(COMP = E.Salary / (E.Age - 18))

where E.Name = "Jones"

QUEL (QUERY LANGUAGE)



SOLD TO
COMPUTER
ASSOCIATES

INGRES

BERKELEY RESEARCH LABS

- ▶ Interdisciplinary groups
- ▶ Lots of collaboration
- ▶ Limited duration
- ▶ Real-world applications
- ▶ Path to commercialization
- ▶ BSD license (usually)
- ▶ Examples: Ingres, Postgres, RISC, Spark, and many more.



**WHAT COMES
POST INGRES?**

POST-GRES.



**STAKE OUT A
CONTROVERSIAL
POSITION, AND
PROVE IT
CORRECT.**

Michael Stonebraker

STONEBRAKER AND ROWE

THE DESIGN OF POSTGRES

THE DESIGN OF POSTGRES, 1986

- ▶ Designed to combine applications & traditional business data
- ▶ Response to Ingres' shortcomings
- ▶ Lots of extensibility
- ▶ Adds a "rules engine"
 - ▶ (predates modern constraints)

THE DESIGN OF POSTGRES

Michael Stonebraker and Lawrence A. Rowe

*Department of Electrical Engineering
and Computer Sciences
University of California
Berkeley, CA 94720*

Abstract

This paper presents the preliminary design of a new database management system, that is the successor to the INGRES relational database system. The new system are to:

provide better support for complex objects,

provide user extendibility for data types, operators and access methods,

provide facilities for active databases (i.e., alerters and triggers) and inferencing, including forward- and backward-chaining,

simplify the DBMS code for crash recovery,

produce a design that can take advantage of optical disks, workstations or multiple tightly-coupled processors, and custom designed VLSI chips, and

make as few changes as possible (preferably none) to the relational model.

This paper describes the query language, programming language interface, system processing strategy, and storage system for the new system.

"DEMOWARE"

1987

POSTGRES 1.0 RELEASED

1989

POSTGRES 2.0, THE FIRST REWRITE

1990

POSTGRES 3.0, THE SECOND REWRITE
(ALSO, EXTENSIBILITY SUPPORT ADDED)

1991

STONEBRAKER FOUNDS ILLUSTRATO
TO COMMERCIALIZE POSTGRES

1992

POSTGRES 4.2

BERKELEY POSTGRES LAB ENDS

1994

POSTGRES IS FINISHED

THE END

... **OR IS IT?**

CHAPTER TWO

ANDREW AND JOLLY



BERKELEY, CALIFORNIA, 1994

**ANDREW AND JOLLY
ARE BORED**

**IT SEEMED LIKE IT WOULD BE
ABOUT AS HARD AS A
UNIVERSITY CLASS PROJECT.**

Jolly Chen



**WE WOULD RECRUIT THE
SMARTEST [STUDENTS] WE
COULD FIND, GIVE THEM
WONDERFUL EQUIPMENT,
AND THEY WOULD
BASICALLY DIE WRITING
CODE FOR US.**

Michael Stonebraker

range of E is EMPLOYEE

retrieve into W

(COMP = E.Salary / (E.Age - 18))

where E.Name = "Jones"

BEFORE: QUEL


```
select (e.salary / (e.age - 18))  
        as comp  
from employee as e  
where e.name = "Jones"
```

AFTER: SQL

UNIVERSITY

POSTGRES 4.2



NOW WITH SQL!

POSTGRES95

SQL SUPPORT?

- ▶ No window functions
- ▶ No foreign keys
- ▶ No inner joins
- ▶ No NULL
- ▶ No constraints
- ▶ No MVCC
- ▶ Not much of anything!



SHIP IT

PROFILE OF AN EARLY ADOPTER

- ▶ Using a SQL database at work
- ▶ Wanted one to run at home
- ▶ Didn't have money to spend
- ▶ Didn't have OS supported by other databases
- ▶ Postgres kinda worked



CHAPTER THREE

POSTGRES IS BORN



BERKELEY, CALIFORNIA, 1996

**ANDREW AND JOLLY
ARE BUSY**



THE INTERNET, 1996

**BRUCE AND MARC
ARE TIRED OF WAITING**

THE BIRTH OF POSTGRESQL

> On Tue, 23 Apr 1996, Jolly Chen wrote:

>

> > I've posted a TODO list on the postgres95 web site

> > (<http://s2k-ftp.CS.Berkeley.EDU:8000/postgres95/www/todo.html>)

> > I've casually sorted the list by priority and I have some editorial

> > comments on some of them.

> >

> > If all the items on the TODO list were completed, postgres95 would be

> > much improved, and would really be a viable replacement for commercial

> > RDBMSs in some settings. Some of the items require quite a bit of work

> > and deep knowledge of postgres95 internals, though. **We would need a few**

> > **contributors with quite a lot of volunteer hours to make this happen**

> > **anytime soon. (A large number of contributors each with only a little**

> > **bit of time to contribute would not be equivalent)**

THE BIRTH OF POSTGRESQL: SCRAPPY STARTS THE REPOSITORY

Date: Mon, 08 Jul 1996 22:12:19 -0400 (EDT)
From: "Marc G. Fournier" <scrappy@ki.net>
Subject: [PG95]: Developers interested in improving PG95?
To: Postgres95 Users <postgres95@oozoo.vnet.net>

Personally, I think that both Jolly and Andrew have done a fantastic job of bringing it to its currently level, but they, like most of the ppl on this list, have full time jobs that sap alot of their time...

...so, unless someone out there has already done this, and unless Jolly/Andrew tell me I can't (guys?)...I'm going to go ahead with what I wanted to do a few months ago...setup a development site similar to what is done with FreeBSD...

First stage will be to get a cvs archive of postgres 1.01 online tonight, with a sup server so that everyone has access to the source code.

If anyone has any patches they wish to submit based off of 1.01, please send them to postgres@ki.net and I'll commit those in as soon as cvs is up and running.

Unless there are any disaggrements with this (or someone else has done this that I missed in mail...sorry if I did...)...I'll send out further data on this as soon as its up and running...

THE BIRTH OF POSTGRESQL: BRUCE THE BUG TRACKER

Date: Mon, 15 Jul 1996 00:28:34 -0400 (EDT)
From: Bruce Momjian <maillist@candle.pha.pa.us>
Subject: [PG95]: Future of Postgres95 development
To: postgres95@oozoo.vnet.net (Postgres95)

I think it would be great if Jolly and Andrew would continue to gather patches from postings to this list, and perhaps help suggest solutions to some of the more complex bugs left in the system (corrupt databases?).

If not, where are we going?

I am willing to bundle up a new release based on patches posted to this group since the release of 1.01.

I am also willing to start gathering a list of known bugs and coordinate efforts to develop patches to fix them.

After these steps, I am willing to post the todo list, and ask for volunteers to start work on them.

NOW ALL WE NEED IS A

NAME



LESS FUNNY IN 1997

POSTGRES95

THE BIRTH OF POSTGRESQL: NAMING

From: "Vadim B. Mikheev" <vadim@sable.krasnoyarsk.su>
Date: Wed, 18 Sep 1996 15:04:22 +0800
Subject: Re: [PG95-DEV] postgres2.0

Marc G. Fournier wrote:

[...]

> How about PostgresV6?

And it's good name. But Ingres & Postgres are different in language (not only), and Postgres & Postgres'95 too, so why not to change prefix now again?

> Please...though...**just not Postgres/SQL**...it just **doesn't slide off the tongue nicely** :(

Vote against too.

THE BIRTH OF POSTGRESQL: NAMING

From: "Marc G. Fournier" <scrappy@ki.net>
Date: Tue, 22 Oct 1996 04:35:34 -0400 (EDT)
Subject: Re: [PG95-DEV] postgres.org

On Tue, 22 Oct 1996, Julian Assange wrote:

>

> Does anyone know what is happening with the delegation of the
> postgres domains? Indirect.com was going to hand them over,
> but as yet has not.

Last week, I sent an email off to the 'admin contact' asking for a status, but have heard nothing... Which is okay.

We had that major discussion on name changes, and it seemed that the majority were in favor of PostgreSQL...so let's do it.

THE BIRTH OF POSTGRESQL: NAMING

From: "Marc G. Fournier" <scrappy@ki.net>
Date: Tue, 22 Oct 1996 04:35:34 -0400 (EDT)
Subject: Re: [PG95-DEV] postgres.org

On Tue, 22 Oct 1996, **Julian Assange** wrote:

>

> Does anyone know what is happening with the delegation of the
> postgres domains? [Indirect.com](http://indirect.com) was going to hand them over,
> but as yet has not.

Last week, I sent an email off to the 'admin contact' asking for a status, but have heard nothing... Which is okay.

We had that major discussion on name changes, and it seemed that the majority were in favor of PostgreSQL...so let's do it.

THE BIRTH OF POSTGRESQL: NAMING

```
commit 9b41da6ce48e3bed6730faa6347a5461175cff83
Author: Bruce Momjian <bruce@momjian.us>
Date:   Wed Dec 11 00:28:15 1996 +0000
```

```
Rename postgres95 to PostgreSQL.
Add comment for SELECT NULL
```

AND SO IT CAME TO BE

"POSTGRESQL"

THE BIRTH OF POSTGRESQL: NAMING

From: Tom Lane <tgl@sss.pgh.pa.us>
Subject: The name of the game
Date: 2006-07-13 03:42:54

It might help to explain that the pronunciation is "post-gres" or "post-gres-cue-ell", not "post-gray-something".

I heard people making this same mistake in presentations at this past weekend's Postgres Anniversary Conference :-(
Arguably, the 1996 decision to call it PostgreSQL instead of reverting to plain Postgres was the single worst mistake this project ever made.

It seems far too late to change now, though.

regards, tom lane

JUST SAY IT LIKE THIS

"POST - GRES"



1997

CHAPTER FOUR

THE GLOBAL DATABASE

THE GLOBAL DATABASE

POSTGRES IS EVERYWHERE




A GLOBAL PROJECT FROM THE BEGINNING: 1997

- ▶ USA
- ▶ Canada
- ▶ Russia
- ▶ Japan
- ▶ Australia
- ▶ Germany

VADIM MIKHEEV

- ▶ The original Russian contributor
 - ▶ Support for MVCC
 - ▶ Write-Ahead-Log
 - ▶ Triggers
 - ▶ Contributed to Russian support
 - ▶ Retired from core team today



ОЛЕГ БАРТУНОВ & ФЕДОР СИГАЕВ

- ▶ Indexes & non-atomic data types
 - ▶ Generalized Search Tree
 - ▶ Generalized Inverted Index
 - ▶ K-Nearest Neighbors Index
 - ▶ Full Text Search
 - ▶ hstore/JSONB
- ▶ "The Mad Russians"

GiST for PostgreSQL



CHAPTER FIVE

COMMERCIALIZATION



THE FIRST SERIOUS
CONTENDER

GREATBRIDGE

\$40,000,000 USD

NED LILLY

VP OF HACKER RELATIONS

TOO EARLY IS THE SAME AS BEING WRONG

GONE IN 16 MONTHS

EDBTM
POSTGRES

THE POSTGRES
COMPANY

ENTERPRISEDB



RUSSIA'S
DEVELOPMENT HOUSE

**POSTGRES
PROFESSIONAL**

COMMERCIAL POSTGRES

FORKS

POSTGRES IS

CLEAN CODE

POSTGRES IS

FREE CODE

```
→ postgres git:(master) X cat COPYRIGHT
PostgreSQL Database Management System
(formerly known as Postgres, then as Postgres95)

Portions Copyright (c) 1996-2016, PostgreSQL Global Development

Portions Copyright (c) 1994, The Regents of the University of C

Permission to use, copy, modify, and distribute this software and
documentation for any purpose, without fee, and without written
is hereby granted, provided that the above copyright notice and
paragraph and the following two paragraphs appear in all copies

IN NO EVENT SHALL THE UNIVERSITY OF CALIFORNIA BE LIABLE TO ANY
DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES
LOST PROFITS, ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS
DOCUMENTATION, EVEN IF THE UNIVERSITY OF CALIFORNIA HAS BEEN AD
POSSIBILITY OF SUCH DAMAGE.

THE UNIVERSITY OF CALIFORNIA SPECIFICALLY DISCLAIMS ANY WARRANTY,
INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANT
AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HERE
ON AN "AS IS" BASIS, AND THE UNIVERSITY OF CALIFORNIA HAS NO OBLI
PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFIC
```

**ANYONE
CAN DO
ANYTHING**

MASSIVELY PARALLEL PROCESSING DATABASES

- ▶ Netezza (7.2)
- ▶ GreenPlum (7.4)
- ▶ RedShift (8.0)
- ▶ AsterData (8.4)
- ▶ CitusData (9.0)

RedShift & GreenPlum are both merging newer code now.

CitusData is now an OSS extension instead of a fork .

CLUSTERED POSTGRESQL

- ▶ Translattice
- ▶ StormDB
- ▶ PostgresForest
- ▶ Postgres-XC
- ▶ Postgres-XL
- ▶ Postgres Pro Enterprise
- ▶ ... and more

PIPELINE DB (9.5)

NEW DATABASES

**FORKS MAKE THE
COMMUNITY STRONG**

CHAPTER SIX

THE MODERN COMMUNITY

NOBODY OWNS POSTGRES

CONSENSUS RULES



[HACKERS]

**MAILING
LIST
CULTURE**

IN SEPTEMBER... OR SO

ANNUAL RELEASES

MONTH-LONG REVIEW PARTIES

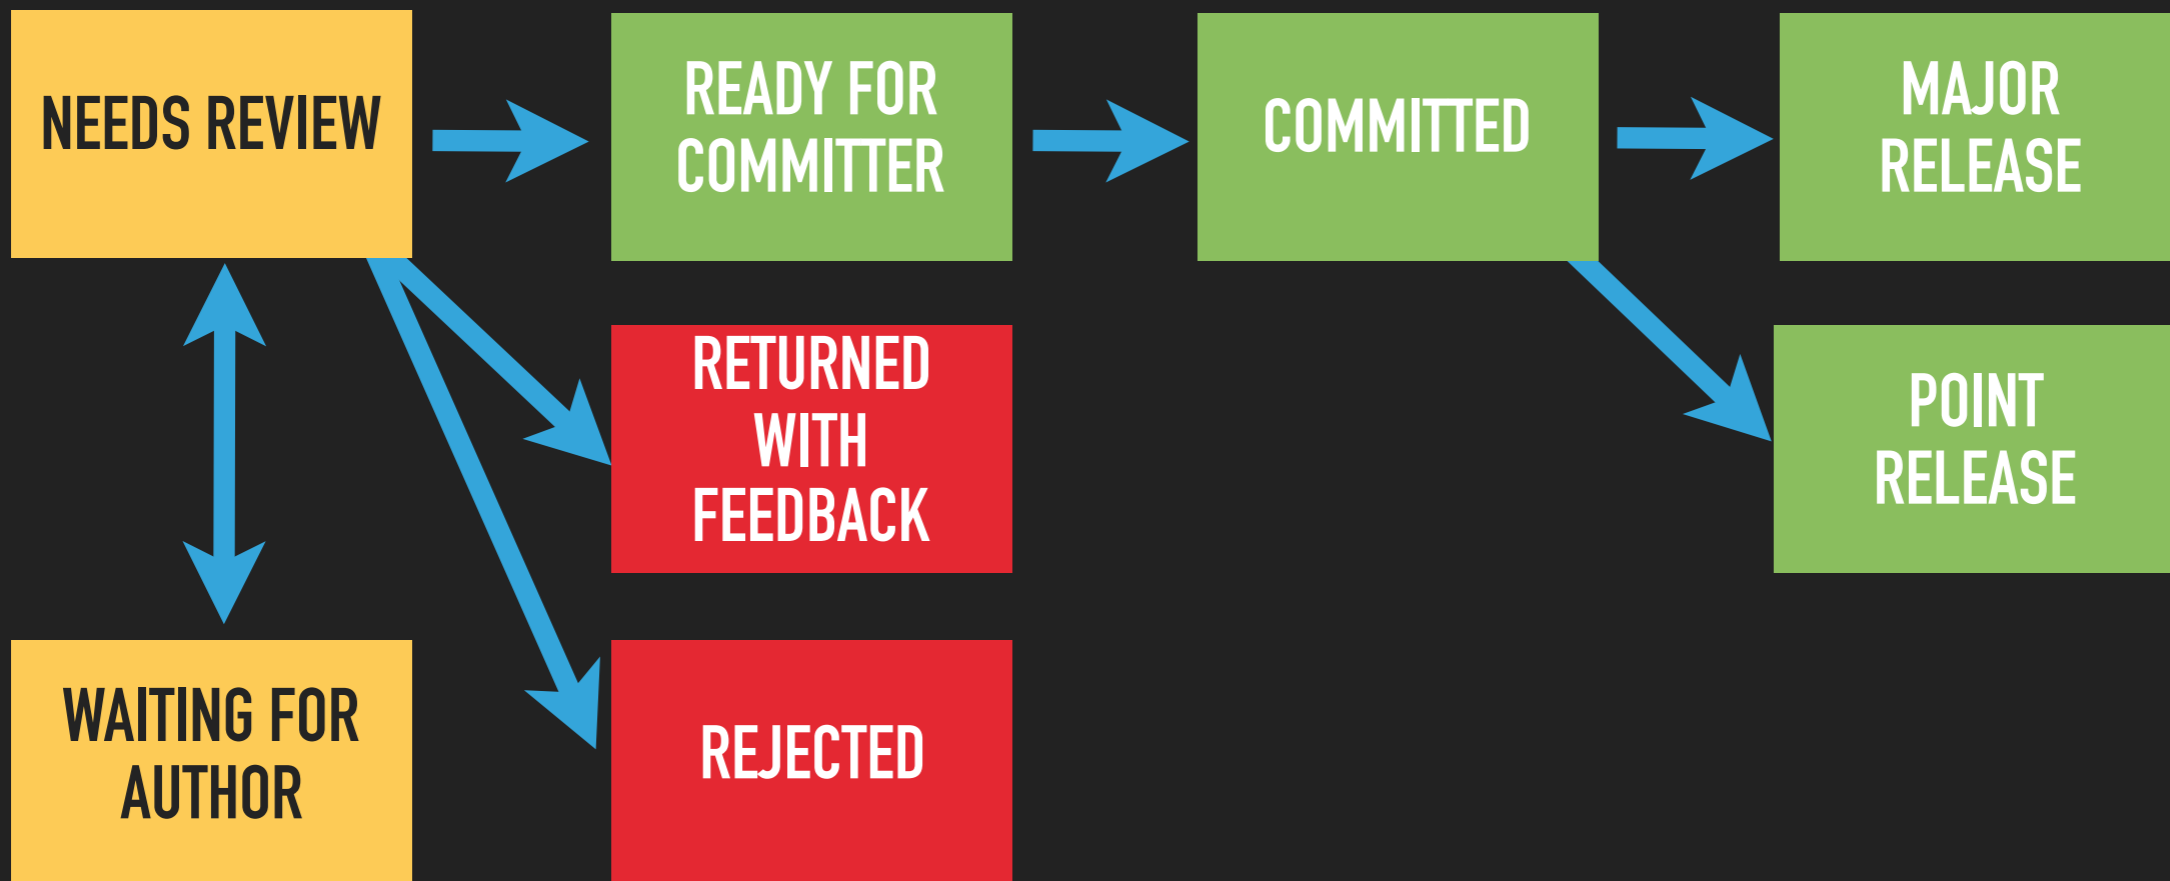
COMMITFESTS

GETTING A PATCH ACCEPTED IN POSTGRESQL

MAILING LIST DISCUSSION




COMMITFEST PROCESS



HOW FEATURES GROW

INCREMENTALISM



DURABILITY

7.1 WAL crash-recovery

8.2 Warm Standby WAL application on standby

9.0 Hot Standby Run queries on standby

9.0 Streaming replication WAL sent continuously

9.1 Synchronous replication WAL can be confirmed by standby

9.4 Logical replication WAL can be represented as logical records

REMOTE DATA ACCESS

7.1 DBLink query remote databases

8.4 SQL/MED
"infrastructure" no user-facing features


9.1 Foreign Data Wrappers very user-facing support

9.2 FOREIGN TABLE support for defining external tables

9.3 postgres_fdw postgres FDW in core

9.6 FOREIGN SCHEMA maintain an external schema

FEATURES GROW



EPILOGUE

**MIRACULOUS
POSTGRES**

**STONEBRAKER
WAS RIGHT.**

**TIMING IS
EVERYTHING.**

PEOPLE MAKE
POSTGRES.

**WE'RE GROWING
FASTER THAN EVER.**

Thanks to
Bruce Momjian,
Oleg Bartunov,
Jolly Chen,
Marc Fournier,
Greg Stark,
Tatsuo Ishii,
Magnus Hagander,
Dave Page,
and everyone else
who contributed interviews,
figures, or reviews.

Inaccuracies are my own!

THANK YOU!

30 YEARS OF POSTGRES