

# *Clarify technical rumours about PostgreSQL*

Dipl.-Inf. Susanne Ebrecht

PostgreSQL Germany

in cooperation with  
PostgreSQL Europe

**FOSDEM 2009**



# *Myth and Rumours*

- Non version upgrade
- You always have to use vacuum
- Not for Windows
- Non full text search
- Non graphical user interface
- Non replication
- Non cluster
- Bad performance
- Not stable
- Complicate usage
- Non enum
- Wrong NULL behaviour
- Non stored procedures
- Non auto increment
- No packages available
- Only for big databases
- Non commercial support
- Arrogant community



# Version Upgrade

Dump



Change  
version



Restore



---

The dream



# *Vacuum*

< 8.1 (2005)



>= 8.1 (2005)



# *PostgreSQL for Windows*

Since 8.0 (released Jan 2005)



# Chapter 12. Full Text Search

<http://www.postgresql.org/docs/current/static/textsearch.html>

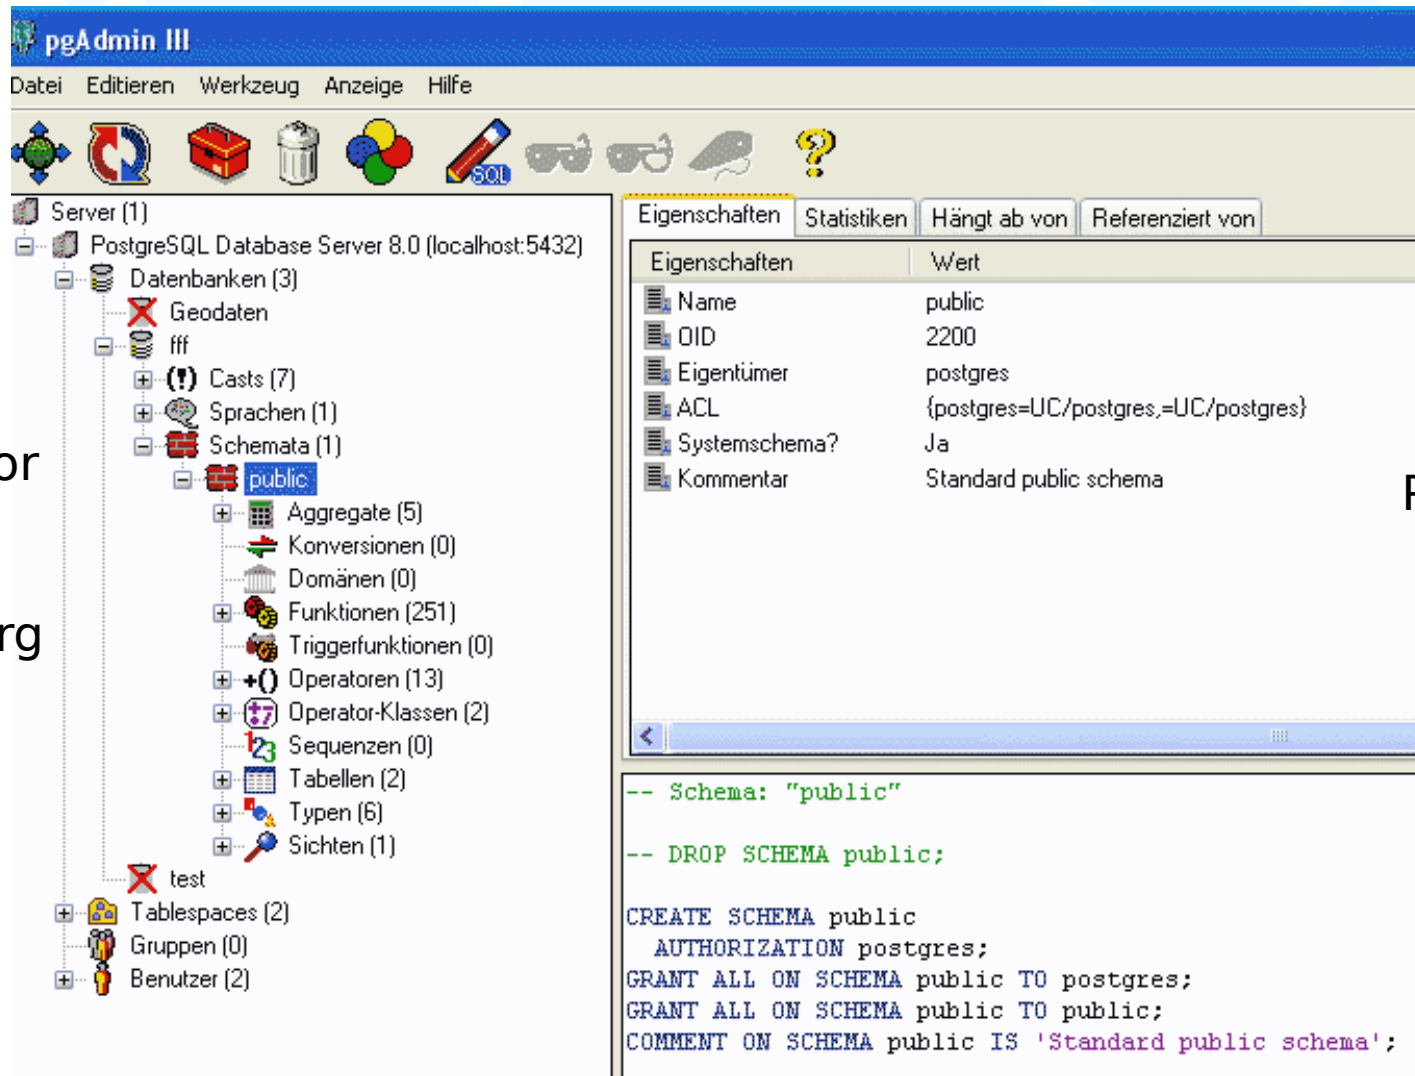
8.3 (2008)



Tsearch2 since 7.4 (2003)

# Graphical User Interface

PgAdmin III (released 2003)



AnySQL  
Maestro

PostgreSQL  
PHP  
Genarator

WaveMaker  
Ajax GUI  
Design Tool

PhpPgAdmin

PgAccess

RHDB  
Administrator

OpenOffice.org

Mergeant

Tora



# *Replication*

SlonyII

NAS

PITR

DRDB

PgPool II

PgCluster

Bucardo

Postgres-R

Fail Over

Load  
Balancing

Remote  
Servers

Data  
Warehousing

Mobile  
Servers



More informations: <http://momjian.us/main/writings/pgsql/replication.pdf>



# Replication and Clustering

| Program   | License | Maturity                    | Replication Method          | Sync         | Connection Pooling | Load Balancing | Query Partitioning |
|-----------|---------|-----------------------------|-----------------------------|--------------|--------------------|----------------|--------------------|
| PGCluster | BSD     | See version details on site | Master-Master               | Synchronous  | No                 | Yes            | No                 |
| pgpool-I  | BSD     | Stable                      | Statement-Based Middleware  | Synchronous  | Yes                | Yes            | No                 |
| pgpool-II | BSD     | Recent release              | Statement-Based Middleware  | Synchronous  | Yes                | Yes            | Yes                |
| slony-I   | BSD     | Stable                      | Master-Slave                | Asynchronous | No                 | No             | No                 |
| Bucardo   | BSD     | Recent release              | Master-Master, Master-Slave | Asynchronous | No                 | No             | No                 |
| Londiste  | BSD     | Recent release              | Master-Slave                | Asynchronous | No                 | No             | No                 |
| Mammoth   | BSD     | Stable                      | Master-Slave                | Asynchronous | No                 | No             | No                 |

[http://wiki.postgresql.org/wiki/Replication%2C\\_Clustering%2C\\_and\\_Connection\\_Pooling](http://wiki.postgresql.org/wiki/Replication%2C_Clustering%2C_and_Connection_Pooling)



# Cluster

## What is cluster?

### SQL:

#### Name

CLUSTER -- cluster a table according to an index

#### Synopsis

CLUSTER tablename [ USING indexname ]

CLUSTER

#### Description

CLUSTER instructs PostgreSQL to cluster the table specified by tablename based on the index specified by indexname. The index must already have been defined on tablename.

...

#### Examples

Cluster the table employees on the basis of its index employees\_ind:

```
CLUSTER employees USING employees_ind;
```

Cluster the employees table using the same index that was used before:

```
CLUSTER employees;
```

Cluster all tables in the database that have previously been clustered:

```
CLUSTER;
```



# Cluster

## What is cluster?

### Database Cluster:

- SQL uses the term catalog cluster.
- PostgreSQL calls it database cluster and it is the initialisation of a database storage area on disk.
- A database cluster is a collection of databases that is managed by a single instance of a running database server.
- You are creating a new cluster by using the command; `initdb`

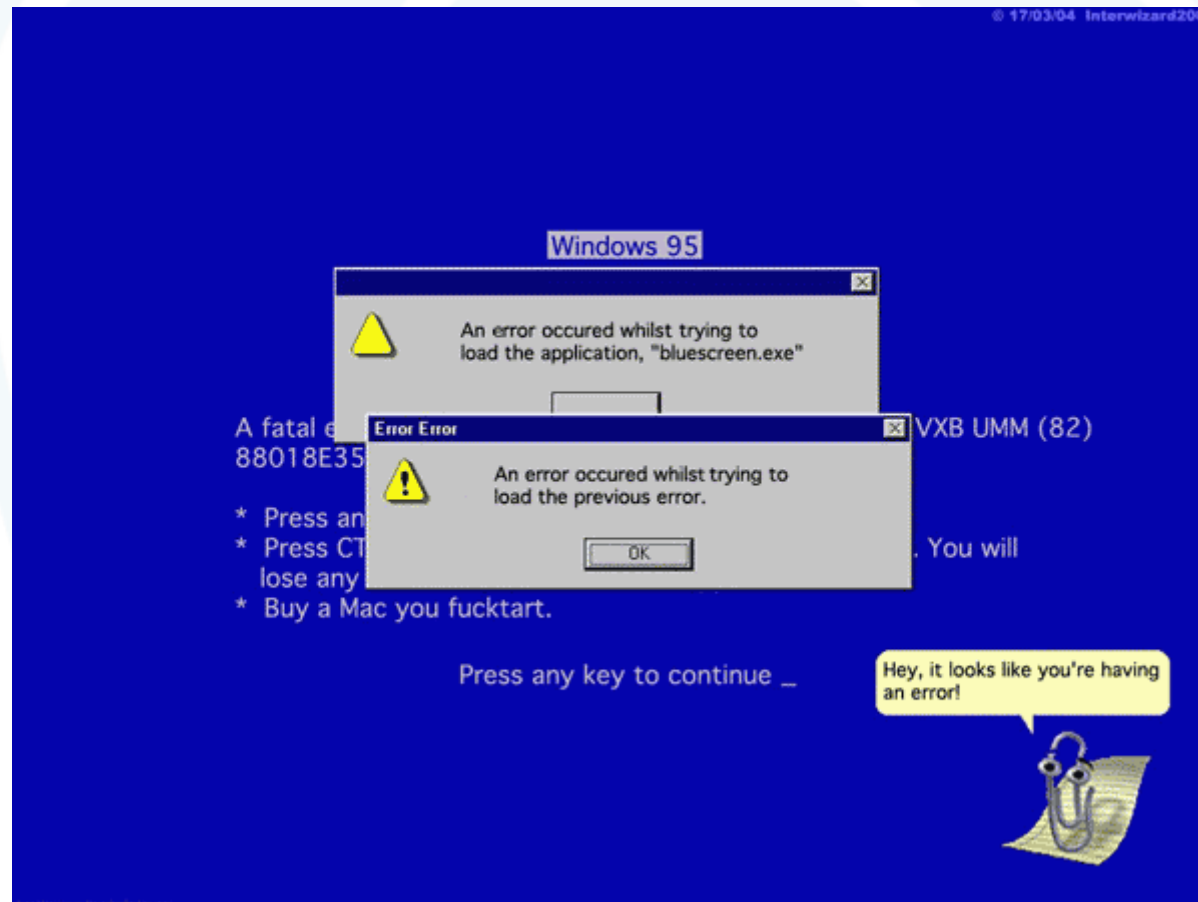


# *Performance*

Which bear has better hunting performance?



# Stableness



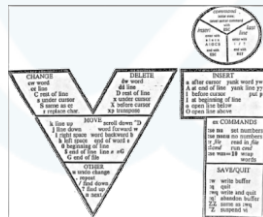
When have you seen a blue screen the last time?

# Usage

$$K = K / q_n ( 1 + ( 100 * ( ( K / (K / q_n ) )_{1/n} - 1 ) ) / 100 )_n$$

Is the formula of compound interest really complicated?  
How fast will you learn these formula when it is your money?

Usage of which editor is more familiar to you?



# *Enum*

## 8.7. Enumerated Types

<http://www.postgresql.org/docs/8.3/static/datatype-enum.html>

Since version 8.3 (2008) PostgreSQL supports ENUM !



# NULL

NULL = 0 ?

NULL = undefined

-1 > 0 < 1 =>

0 = defined

undefined != defined

**NULL != 0**

qed

NULL = " ?

NULL = undefined

" = empty String =>

empty is a clear definition =>

empty String = defined

undefined != defined

**NULL != "**

qed

# *Stored Procedures*

According to SQL Standard:

Procedures are functions with non return value.

PostgreSQL:

Ok, why should we make it more complex?

Just let functions have the possibility to return void.

Summary:

In PostgreSQL stored procedures exist.

Stored procedures are functions returning void.

# *Auto Increment*

## What is auto increment?

Meaning of auto increment is a sequence used in nextval() function. The sequence has a positive number as start value and a positive increment. Default for start and increment is 1.

PostgreSQL sequences can do lots of more. You are able to count backward; you can use negative values; you can define minimum and maximum value; it is your decision if you want to wrap around if minimum or maximum value is reached or if you won't wrap.



# *Packages*



<http://www.postgresql.org/download/>

## **Binary packages**

Pre-built binary packages are available for a number of different operating systems:

**FreeBSD · Linux · MacOS X · Solaris · Window**



# *Big or small databases?*

The baby cat fits into the adult cat basket as well



# *Support?*

IRC (Freenode):

#postgresql

#postgresql-de

#postgresql.fr

#postgresql-it

...

Mailing lists:

<http://www.postgresql.org/community/lists/>

Forums

Commercial support:

[http://www.postgresql.org/support/professional\\_support](http://www.postgresql.org/support/professional_support)

Author recommendations:

IRC or Mailing list

Companies:

Sun Microsystems (PostgreSQL on Solaris)

credativ (placed in Germany, acting world wide)



# *Community character*

Look here ...  
Visit our booth ...  
Chat with us!

Do you really think that we are arrogant?



# *Closing words*

Thanks for listening

