

Success is 99% Persistence

Simon Riggs
CTO, 2ndQuadrant



Simon Riggs

- CTO, 2ndQuadrant
- 25 years working with Business Databases
- 10 years with PostgreSQL
- 8 years with Oracle and DB2
- 7 years with Teradata
- 2 years in academia – paper to Royal Society

What does it take to Succeed?

- Success
 - The achievement of something desired, planned or attempted
 - The accomplishment of one's goals
- Persistence
 - The quality that allows someone to continue doing something even though it is difficult or opposed, often for extended periods of time
 - In computer science, persistence refers to the characteristic of state that outlives the process that created it

Realisations

- To succeed you must have a goal
- You move from **not-success** to **success**
- You must have a plan to take you there
- To persist without a plan is merely insane, as Einstein said: Insanity is doing the same thing again and again and expecting a different outcome
- Plans always go wrong, so even with a plan you will need great persistence

Case Study: Egyptian Pyramids


- Technical Challenges
- Workforce Recruitment & Motivation
- Vision
- Funding Model



Solutions: Egyptian Pyramids

- Technical Challenges
 - Prototyping
 - Incremental Build
- Workforce Recruitment & Motivation: Zeal or whips?
- Vision: Radical design concept,
Don't be seen dead without it
- Funding Model: Pharaoh

- Persistence required
across decades





Case Study: Robust Postgres

- Technical Challenges
- Workforce Recruitment & Motivation
- Vision
- Funding Model

Replication

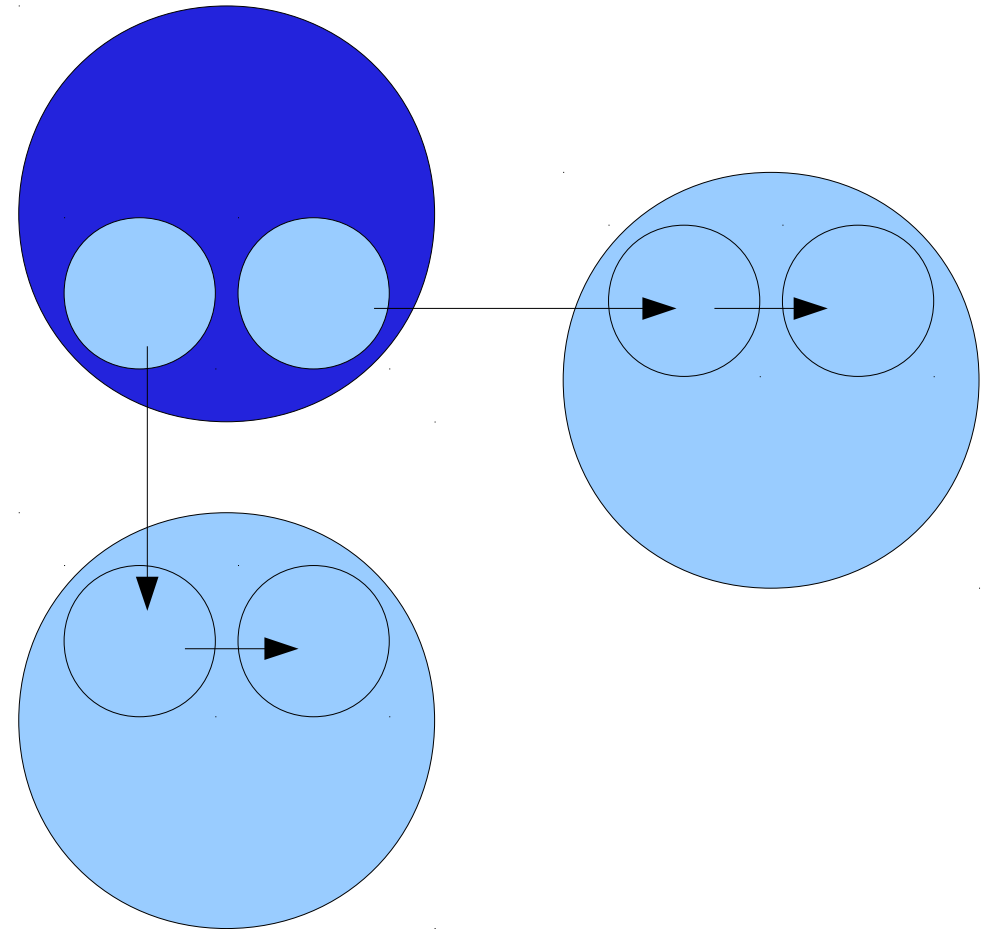
rep·li·ca·tion /,repli'kāSHən/

Noun: The action of copying or reproducing something

- PostgreSQL Databases
 - **Mechanism for maintaining a live copy of a database for purpose of High Availability and Disaster Recovery**
 - Sending data from one database to another, to allow additional applications to access and/or enhance data


Efficient and Concurrent

- Each Receiver node has a dedicated process on Sender node
- Each Receiver node has a dedicated process for communications and another for applying changes



Cascading Replication

- Receivers can also send data onto other Receivers
- Synchronous replication with transaction controlled robustness



Related Developments

- 2004
- 8.0 – Point in Time Recovery
 - 8.2 – Restartable Recovery, Log Switching
 - 8.3 – Full page optimization, pg_standby
 - 8.4 – BgWriter during Recovery
 - 9.0 – Streaming Replication, Hot Standby
 - 9.1 – Synchronous Replication
 - 9.2 – Cascading Replication
- 2013
- 9.3 – Fast Failover, Timeline following



Solutions: Robust Postgres

- Vision: Zero Data Loss
- Technical Challenges
 - Incremental Build
- Workforce Recruitment & Motivation: Zeal
- Funding Model
 - Incremental steps, Viral approach

- Persistence required across a decade

The Importance of Persistence

In computer science, persistence refers to the characteristic of state that outlives the process that created it

- Databases hold a record of current state
- Databases provide transactional controls that allow us to record changes in state accurately and allow concurrent access to that information
 - Bookkeeping systems
 - Booking systems
 - Resource control systems
- Single Version of the Truth

Analytical Databases

- Databases also record transactional history that allows us to analyse what happened
 - Business Intelligence
 - Data Warehousing
 - Just in Time Fulfilment
 - Efficient Manufacturing & Distribution
- Full, detailed data at event level
 - Make business processes visible => improvable
 - “Big Data”
- Clear, agreed meanings for stored data



Successful Businesses use Well Planned Databases

- British Airways
- Wal*Mart
- Bank of America
- Oxfam
- Vodafone
- Navionics
- BBC
- McAfee
- AbnAmro
- Skype
- TomTom
- Open University
- NTT



Successful Technologies

- Relational Databases
 - SQL
 - MVCC
 - Data Mining
-
- Databases are an important part of success in business. They aren't always easy and they take time, money and persistence to make them work




Goals

- To produce the World's Most Advanced General Purpose Database as an enabling technology for ethical business, science, government, health and happiness
- 10 years in...
- 10+ years to go



2ndQuadrant

Professional PostgreSQL



- Professional Development
- 24x7 Production Support and RemoteDBA
- Training, Advice and Doing Stuff For You

- info@2ndQuadrant.com