How we made Greenplum Open Source

Andreas Scherbaum

Pivotal

Andreas Scherbaum

- Works with databases since ~1997, with PostgreSQL since ~1998
- Founding member of PGEU
- Board of Directors: PGEU, Orga team for pgconf.[eu|de], FOSDEM
- PostgreSQL Regional Contact for Germany
- Ran my own company around PostgreSQL for 7+ years
- Joined EMC in 2011, then Pivotal, then EMC, then Pivotal, working on PostgreSQL and Greenplum projects
- Run PostgreSQL Meetups in Palo Alto, CA

Pivotal

Agenda

Agenda

- What is Greenplum?
- History of Greenplum (aligned to PostgreSQL)
- Why Open Source?
- Challenges
- Lessons Learned
- Q & A

Pivotal

What is Greenplum?

Time for Marketing …

Pivotal

What is Greenplum Database?

- One of these PostgreSQL forks
- Diverged away a long time ago (2007)
  - And now wants to merge again (2015 - ???)

What is Greenplum Database?

- Massive-parallel, shared-nothing database
- Optimized for OLAP and Analytics workloads

Why another database?

- PostgreSQL is good for OLTP
  - It does not scale well, it does not scale across dozens of machines
- OLAP systems today scale to 100s of TBs, or even PBs
- Huge amounts of data must be processed
  - Contrary to “hot spots” in OLTP
- Trend from “Store-everything-Data-Warehouse” to “Analytics-Database-on-everything”

Pivotal

Google Trends – Data Warehouse

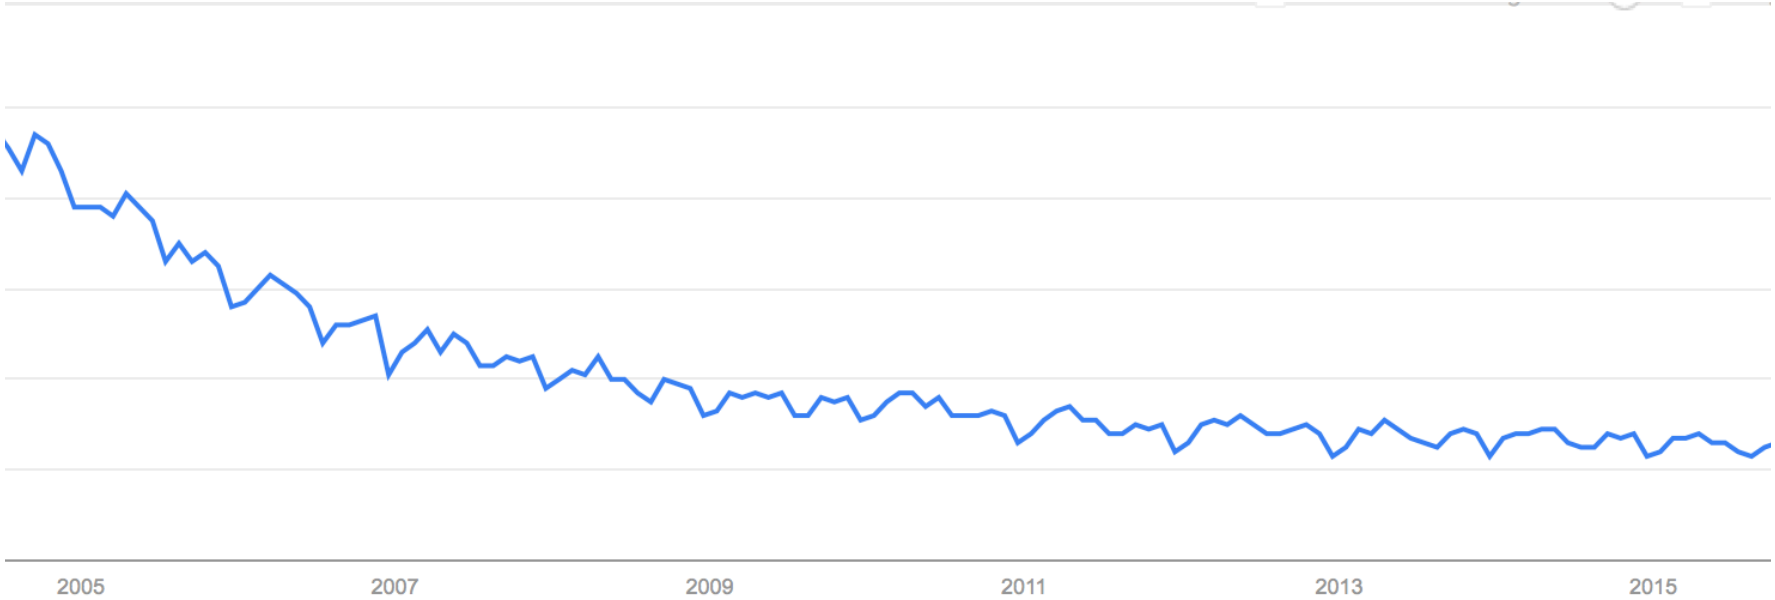

Google Trends – Analytics Database

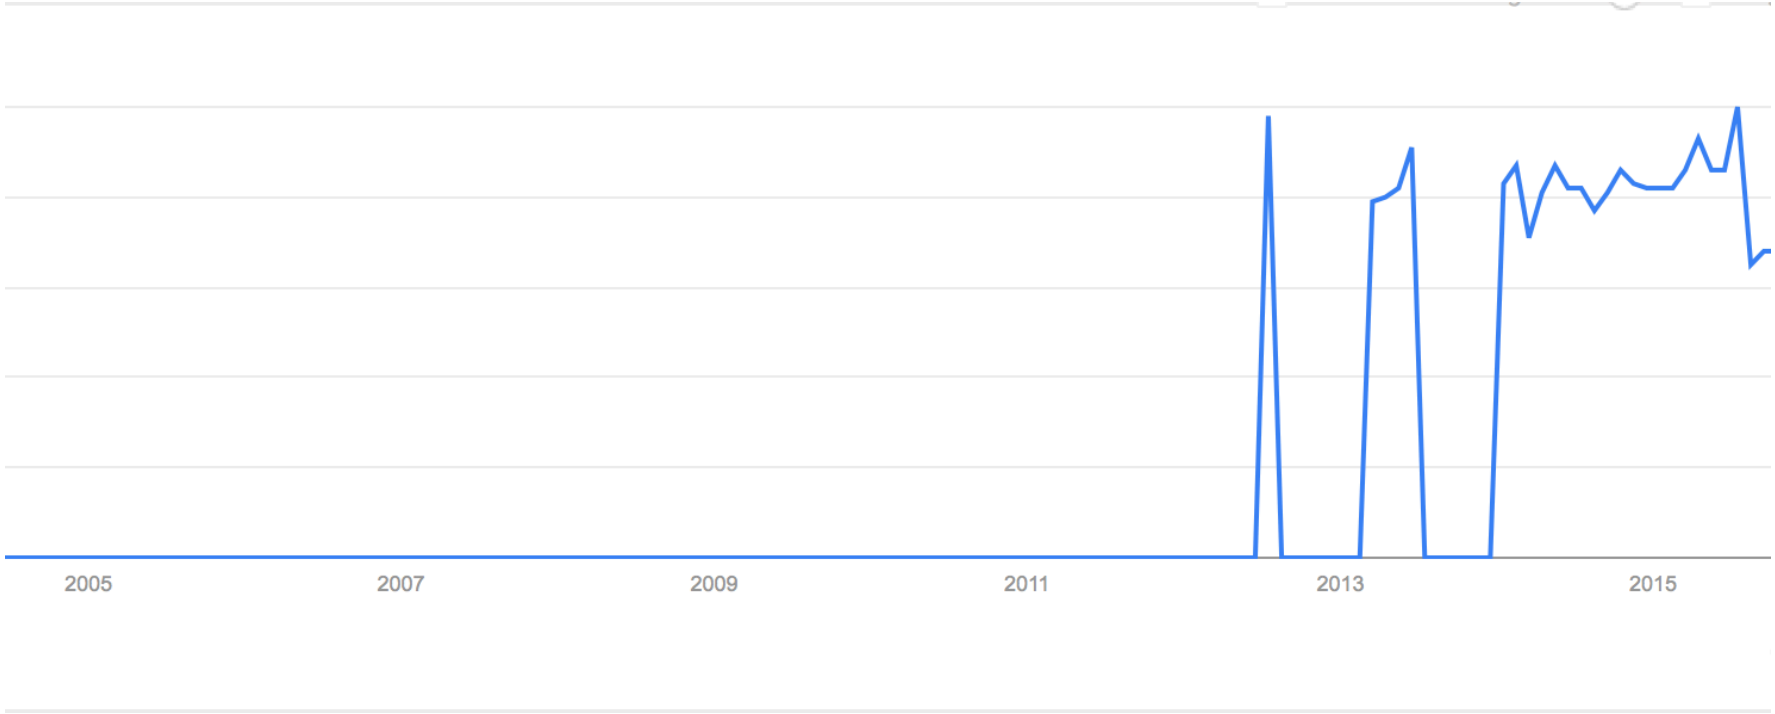

Why another database?

- Many vendors in the market
  - Teradata (founded back in 1979)
  - Greenplum (2003)
  - Exadata (2008)
  - Netezza (1999)
  - …

What is Greenplum Database?

- Massive-parallel, shared-nothing database

- What?

Massive parallel

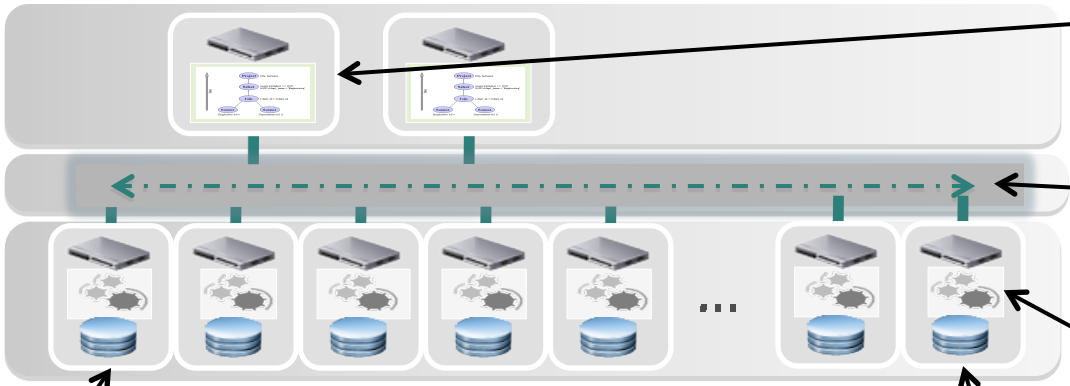

Master server controls everything

Interconnect: Communication between everyone

Every segment server holds n segment databases

1 (2) Segment Server

n Segment Server

Shared nothing

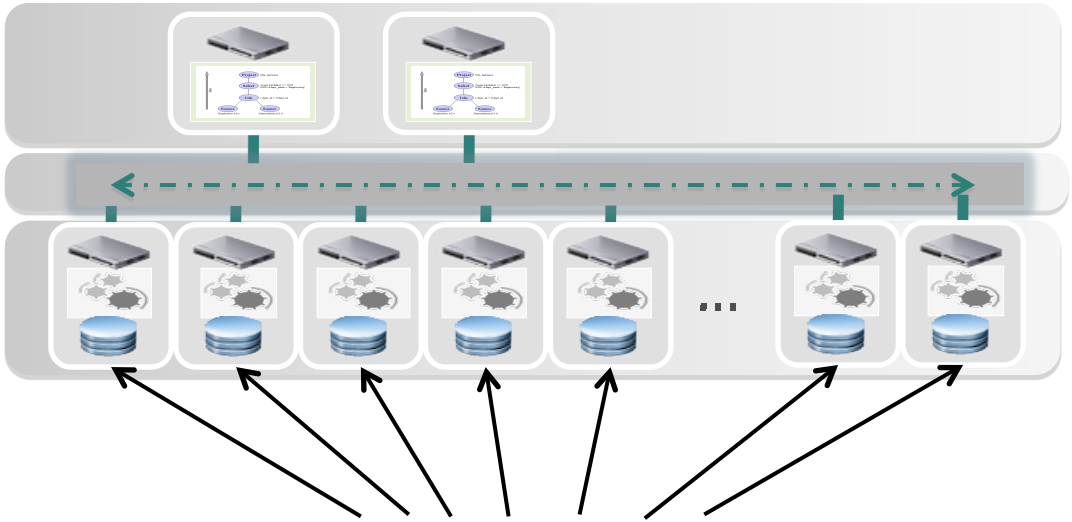

Separate servers:
Share no hardware or storage

Where is my data?

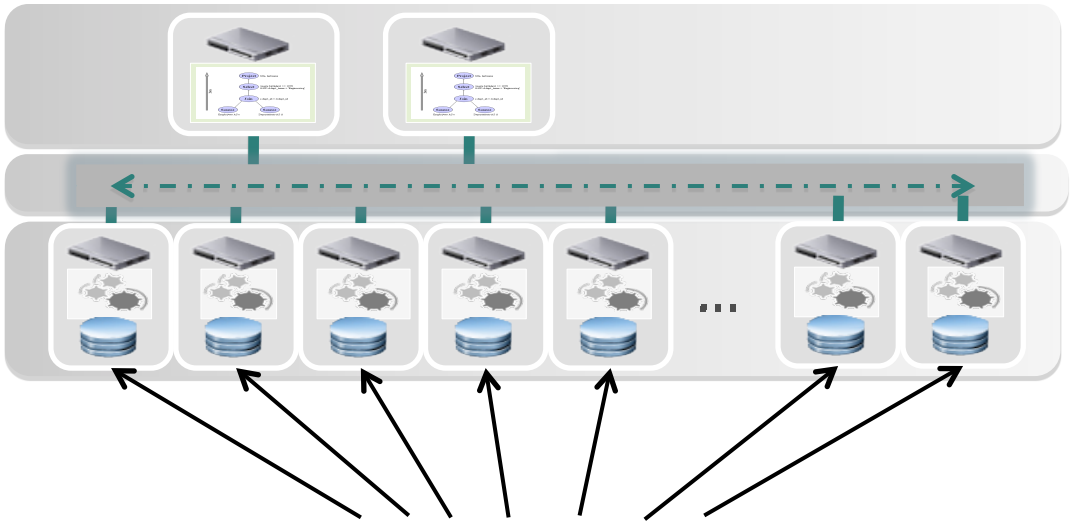

Data is spread across the segment databases
(hopefully evenly distributed)

Segment database with biggest result set will need most time for processing query

Which parts executes my query?

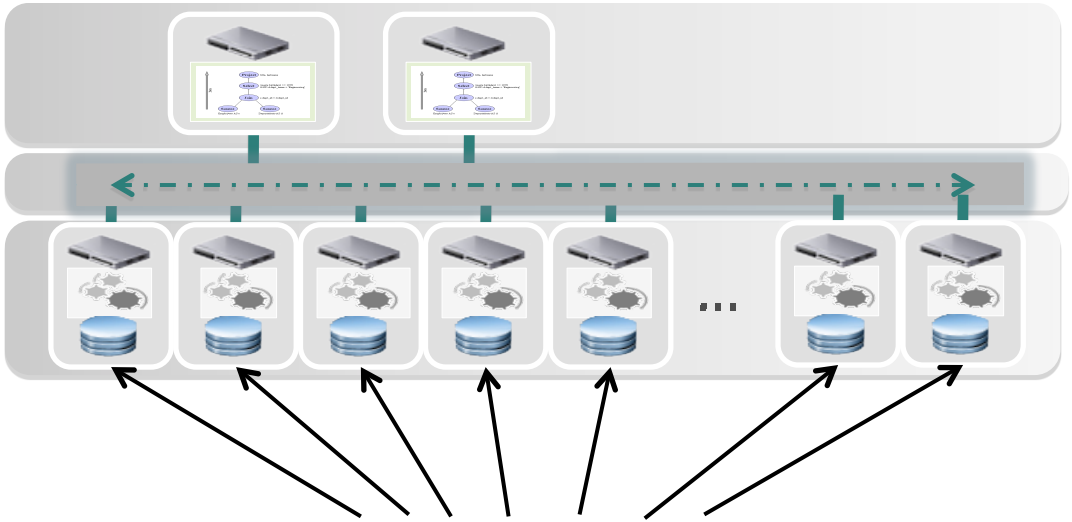

All segment databases, in parallel, on their own potion of the data

100% CPU usage guaranteed[1]

([1]: conditions apply)

But what about joins?

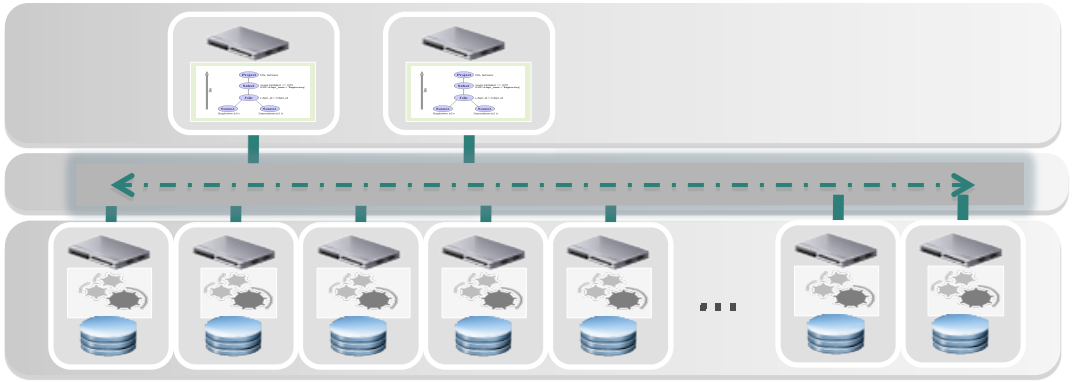

Optimal case:
Data is co-located on the same segment databases

No network traffic

And without co-location?

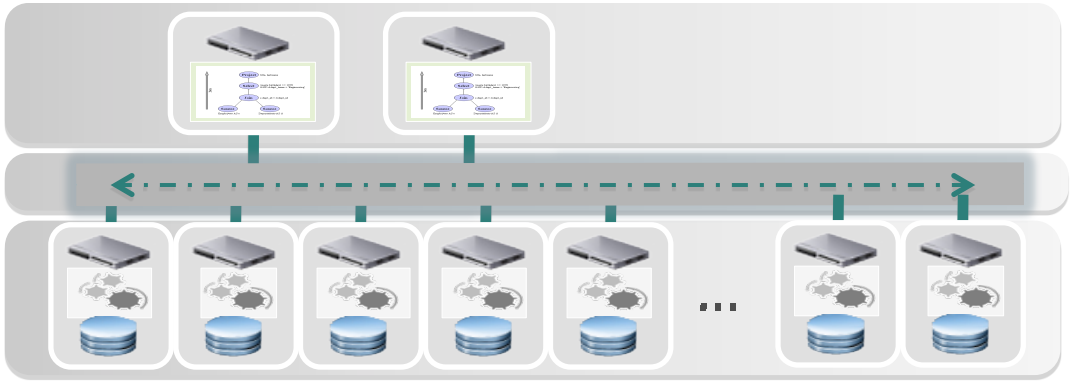

Data is broadcasted or redistributed over the Interconnect

Redistribution:
Small tables

Broadcast:
Big tables

Greenplum is handling the Join

Expansion adds new resources

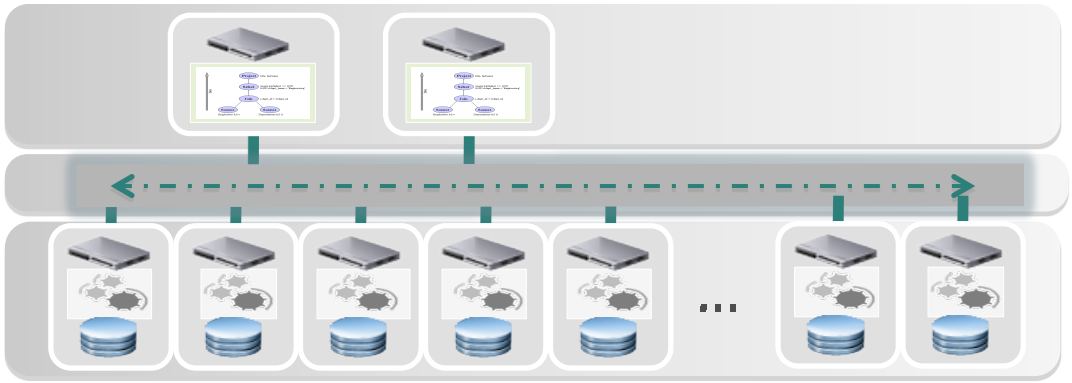

More hardware:
= more CPU cores
= more RAM
= more disk space
= more I/O

Expansion = new hardware

What else?

- Number of Big Data features

Polymorphic storage

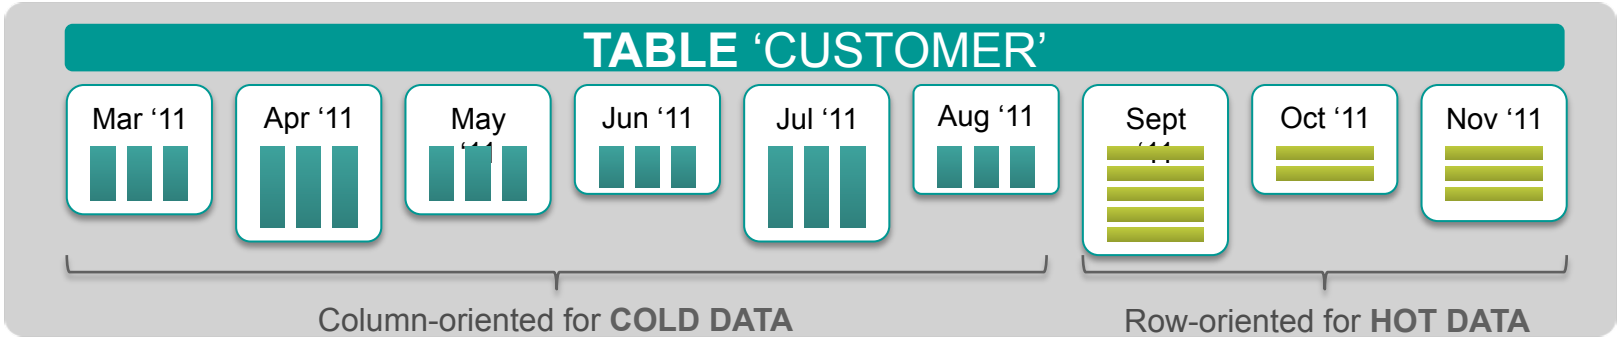

Plus compression

Partitioning (including proper SQL Syntax)

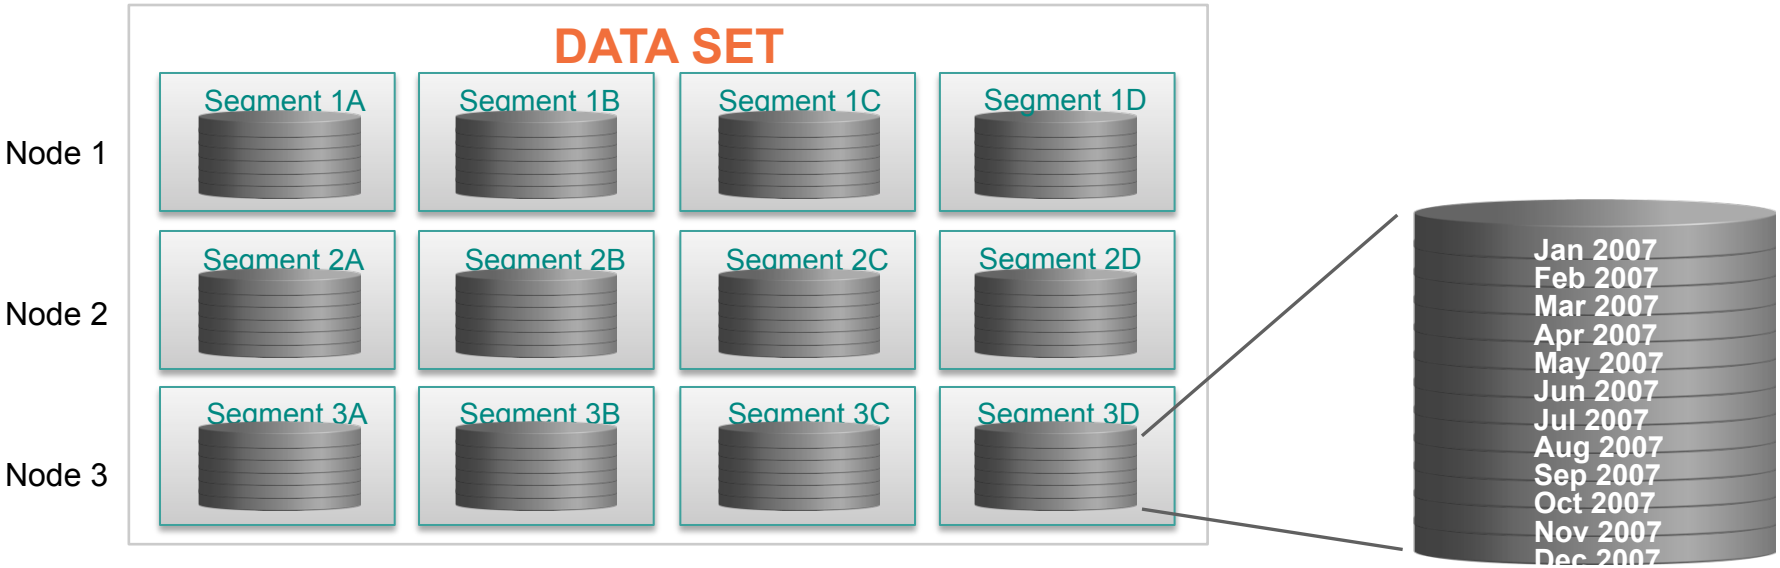

Parallel Data Loading

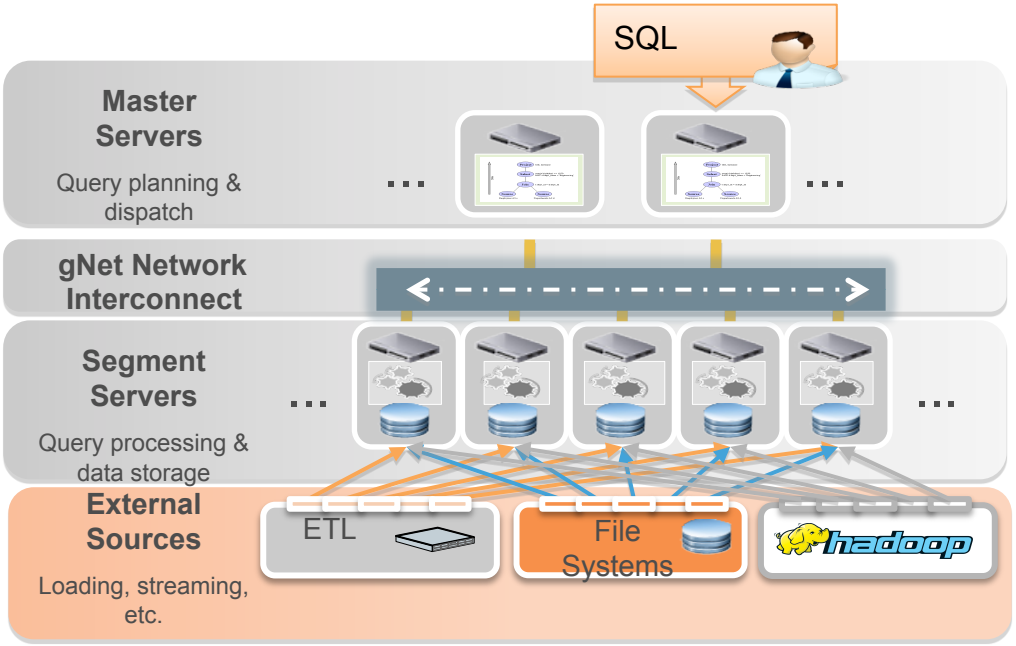

History of Greenplum

Aligned to PostgreSQL timeline

Pivotal

History of Greenplum

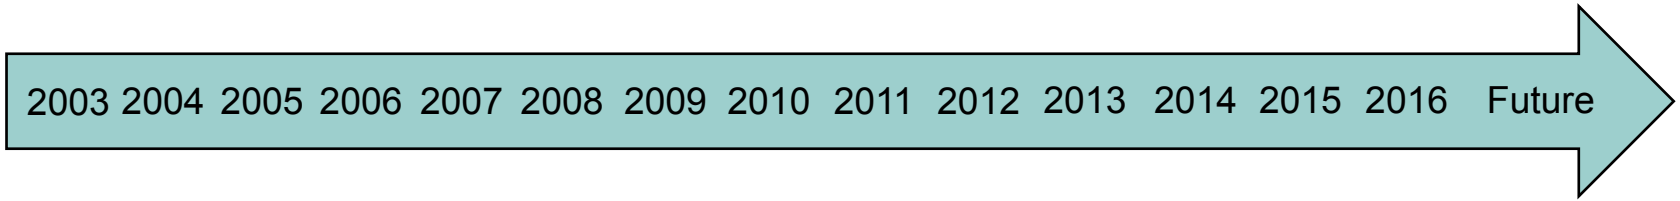

Pivotal

History of Greenplum - 2003

September 2003:
Greenplum founded

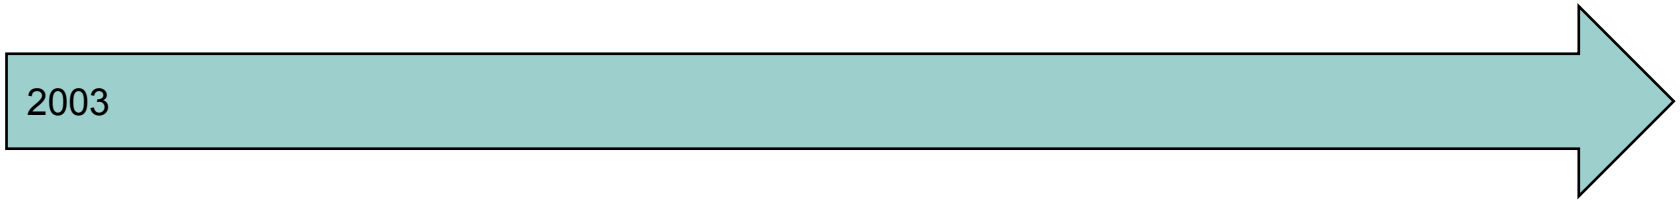

November 2003:
PostgreSQL 7.4
released

Pivotal

Greenplum founded

- Merger of two companies:
  - Didera (DC) + Metapa (NY)
  - Luke Lonergar + Scott Yara
  - Company moved to San Mateo
- Chief Architect: Chuck McDevitt
  - Employee #20 at Teradata
- Products like “BizGres”
  - Discussion: Single node OS, multi node commercial

Pivotal

History of Greenplum – 2004, 2005

…

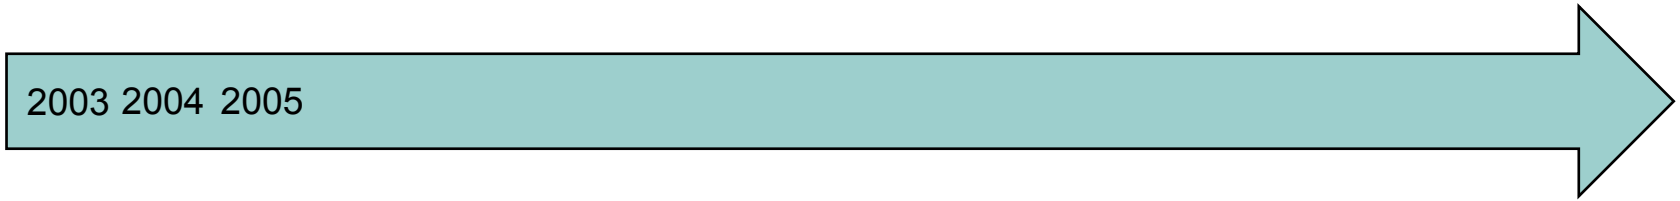

January 2005:
PostgreSQL 8.0
released

November 2005:
PostgreSQL 8.1
released

Pivotal

History of Greenplum - 2006

…

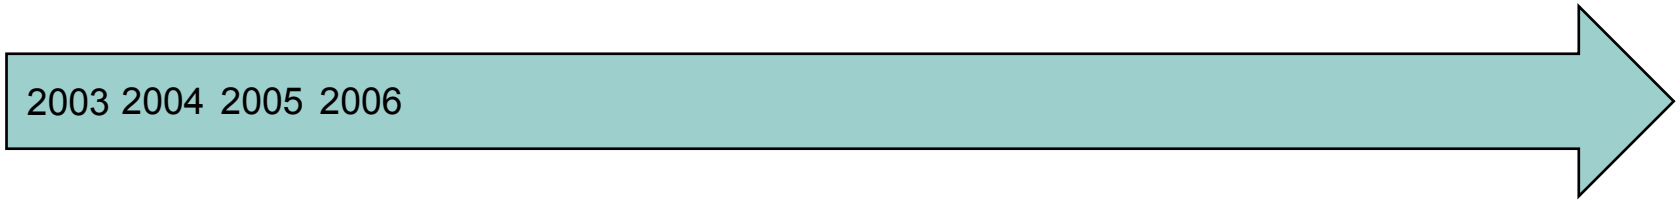

December 2006:
PostgreSQL 8.2
released

History of Greenplum - 2007

2007:
Greenplum 3.0
released

Stopped merging
with PostgreSQL

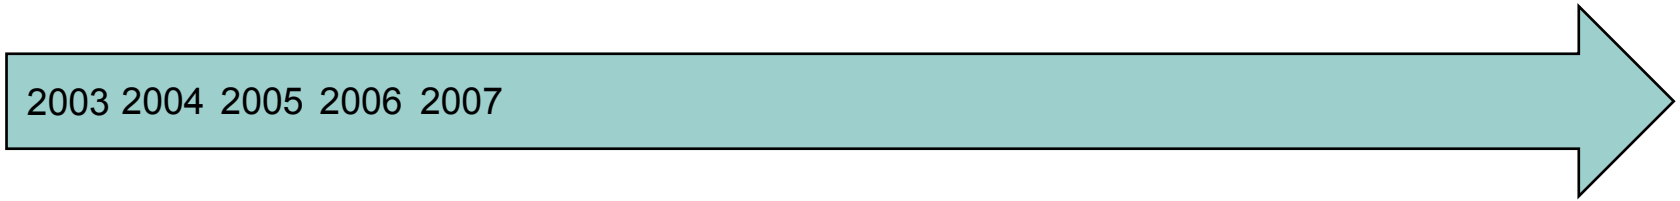

Pivotal

History of Greenplum - 2008

2008:
Greenplum 3.1
released

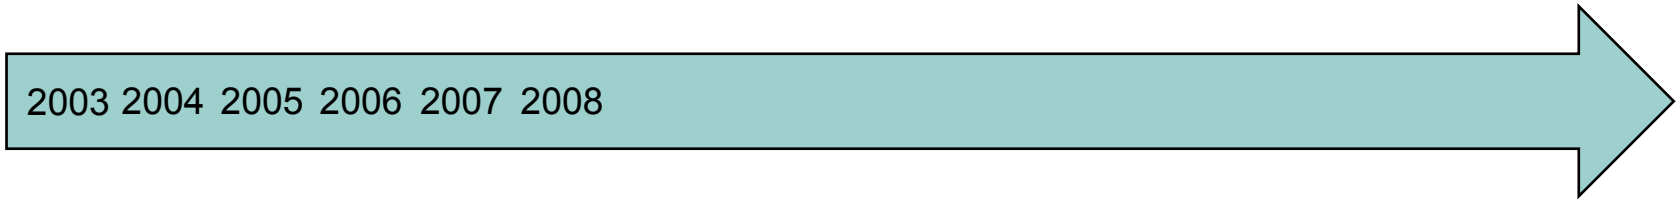

February 2008:
PostgreSQL 8.3
released

Greenplum 3.x

- Logical replication
  - PostgreSQL got replication with v9.0, 3 years later
- Cooperation with Sun
  - GPDB on Sun hardware, running on Solaris
- Stopped merging with PostgreSQL
  - GPDB up to today shows “version 8.2.14”

Pivotal

History of Greenplum - 2009

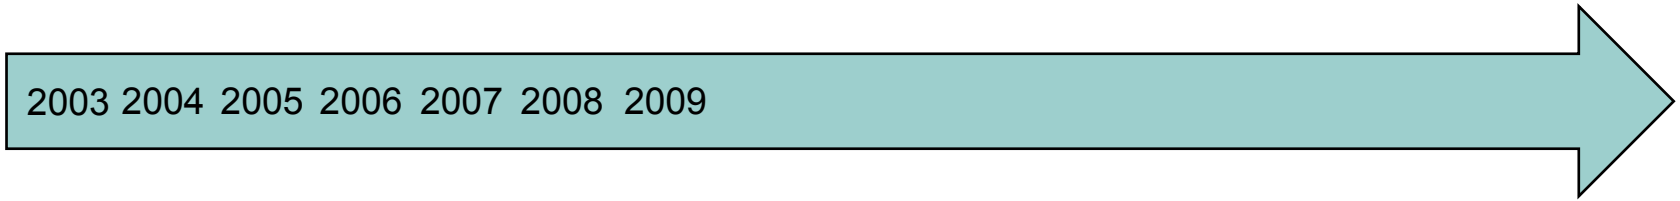

July 2009:
PostgreSQL 8.4
released

History of Greenplum - 2010

May 2010:
Greenplum 4.0
released

July 2010:
Greenplum
acquired by EMC

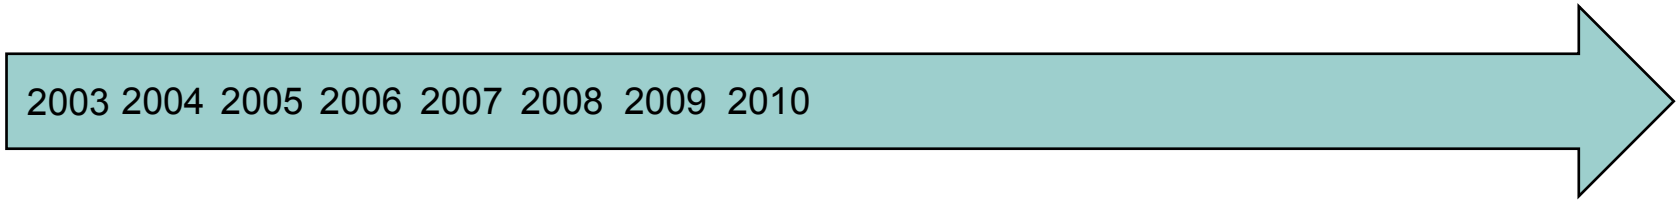

September 2010:
PostgreSQL 9.0
released

Pivotal

Acquisition by EMC

- Already teams in: New Zealand, Australia, China, Silicon Valley, Virginia, Israel, India
- Data Computing Appliance (DCA) v1 released

Pivotal

Greenplum 4.0

- Filesystem based replication
  - Every writing file I/O is replicated from primary to mirror
  - Before that: single node failure made the cluster read only
- Data Science Team

History of Greenplum - 2011

Mar 2011: Greenplum 4.1 released

November 2011: Greenplum 4.2 released

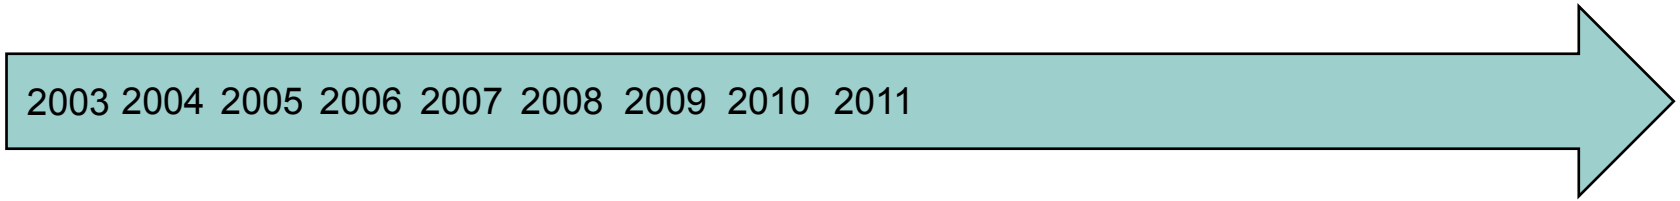

September 2011: PostgreSQL 9.1 released

Pivotal

2011

- MADlib
  - Mad Skills paper “MAD Skills: New Analysis Practices for Big Data”

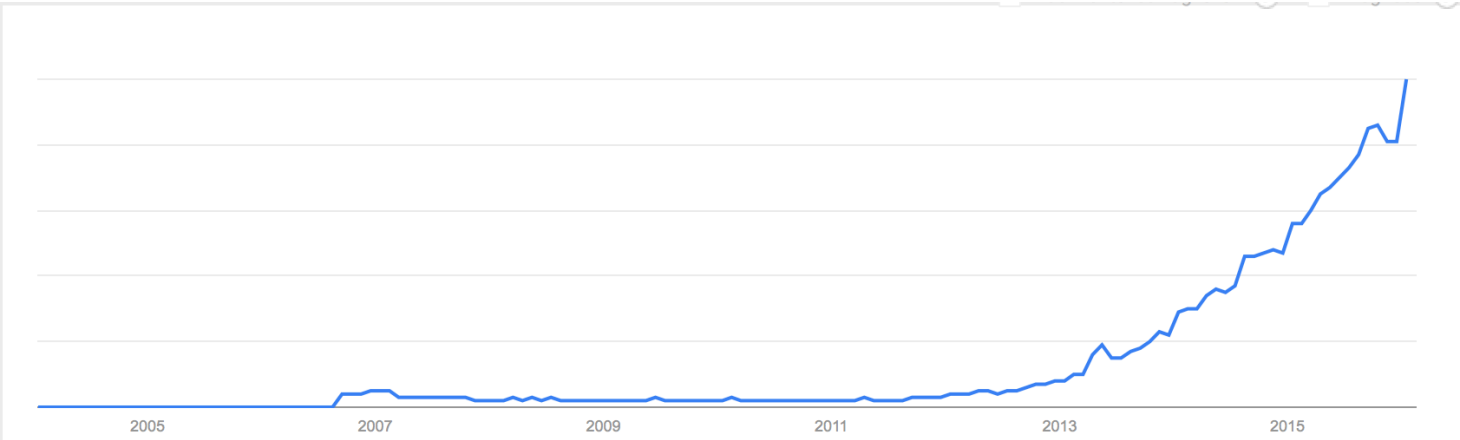

Google Trends: Data Science

2011

- Greenplum Hadoop distribution
- Interconnect switches to UDP
  - TCP does not scale
  - Packet verification on top of UDP

Pivotal

History of Greenplum - 2012

November 2012:
Greenplum 4.2.3
released

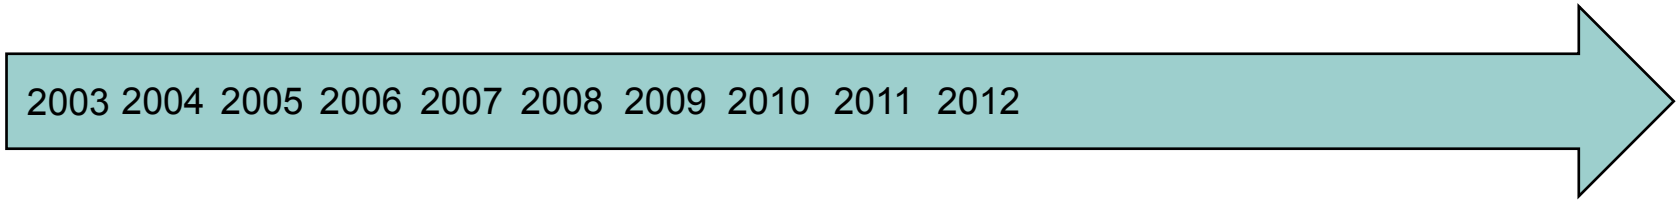

September 2012:
PostgreSQL 9.2
released

Pivotal

History of Greenplum - 2013

February 2013:
Greenplum 4.2.4
released

April 2013:
Greenplum 4.2.5
released

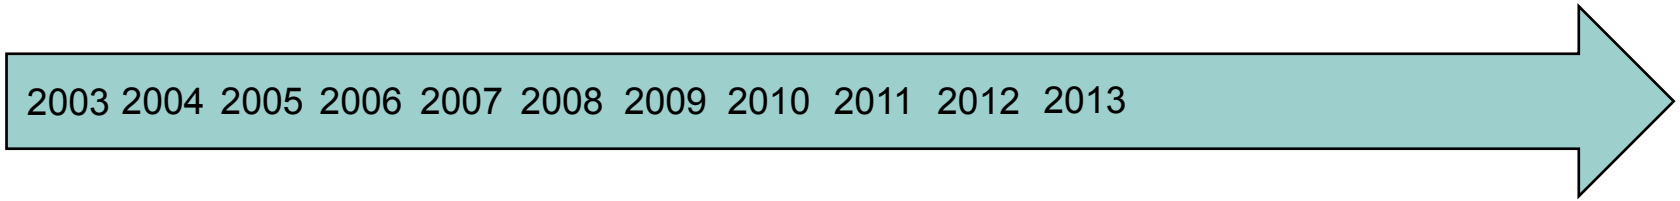

September 2013:
PostgreSQL 9.3
released

Pivotal

2013

- April 2013: EMC & VMware spin off Pivotal

History of Greenplum - 2014

January 2014:
Greenplum 4.3
released

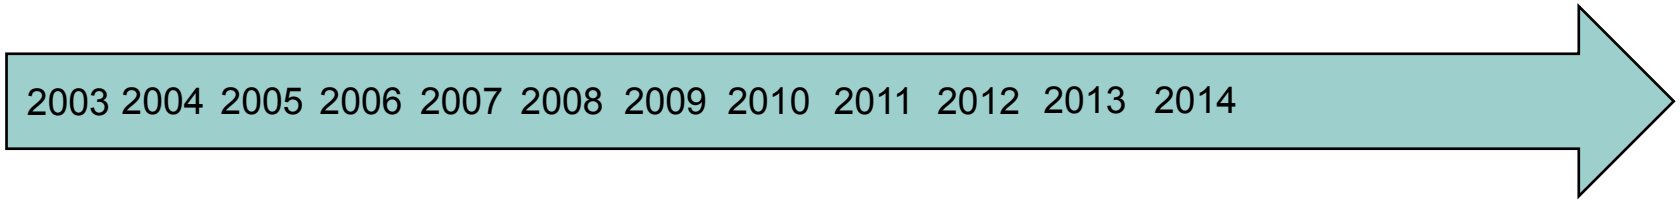

December 2014:
PostgreSQL 9.4
released

Pivotal

2014

- Append-only tables -> Append-optimized tables
  - Allows update in append-only tables
- WAL Replication backported from PostgreSQL
  - For Master -> Slave replication
  - Primary -> Mirror still uses file replication

Pivotal

History of Greenplum - 2015

Open Source

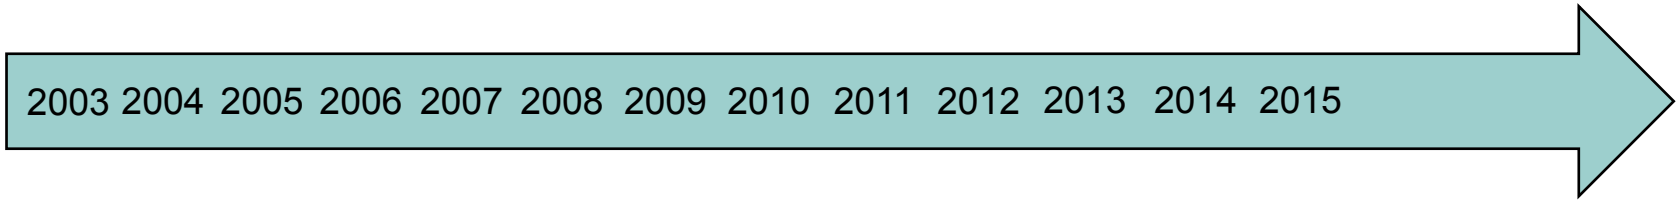

…

2015

- February 17: Pivotal announces that it will Open Source most of it’s products

Pivotal

San Francisco, February 17, 2015 – Pivotal, a

Summary

Pivotal intends to open source core components of its big data products: Pivotal HAWQ®, Pivotal Greenplum® Database and Pivotal GemFire®.

Pivotal Big Data Suite to offer enterprises greater data agility and speedier cloud deployments.

Single subscription for big data deployments for application developers and data analysts.

Pivotal Big Data Suite to include new entitlements for Pivotal Cloud Foundry and new application services.

Strategic alliances to provide a unified and open approach for Hadoop, removing lock-in and providing a more stable and predictable Hadoop distribution.

At 11:00am PST today, Pivotal will host a live streaming event to expand on today's news. To join, visit http://bit.ly/1vR629D.

And after the announcements?

- Silence for several months …

Pivotal

And after the announcements?

- Silence for several months …
- But in the background, a lot happened
  - Number of PostgreSQL contributors join Pivotal:
  - Heikki Linnakangas, Daniel Gustafsson, Dave Cramer, Atri Sharma

Pivotal

And after the announcements?

- Silence for several months …
- But in the background, a lot happened
  - Number of PostgreSQL contributors join Pivotal:
  - Heikki Linnakangas, Daniel Gustafsson, Dave Cramer, Atri Sharma
- June 2015: Apache Geode (Gemfire)

Pivotal

And after the announcements?

- Silence for several months …
- But in the background, a lot happened
  - Number of PostgreSQL contributors join Pivotal:
  - Heikki Linnakangas, Daniel Gustafsson, Dave Cramer, Atri Sharma
- June 2015: Apache Geode (Gemfire)
- October 2015 (pgconf.eu): Greenplum Database
- January 2016: Apache HAWQ (SQL on Hadoop)

Pivotal

Greenplum Open Source

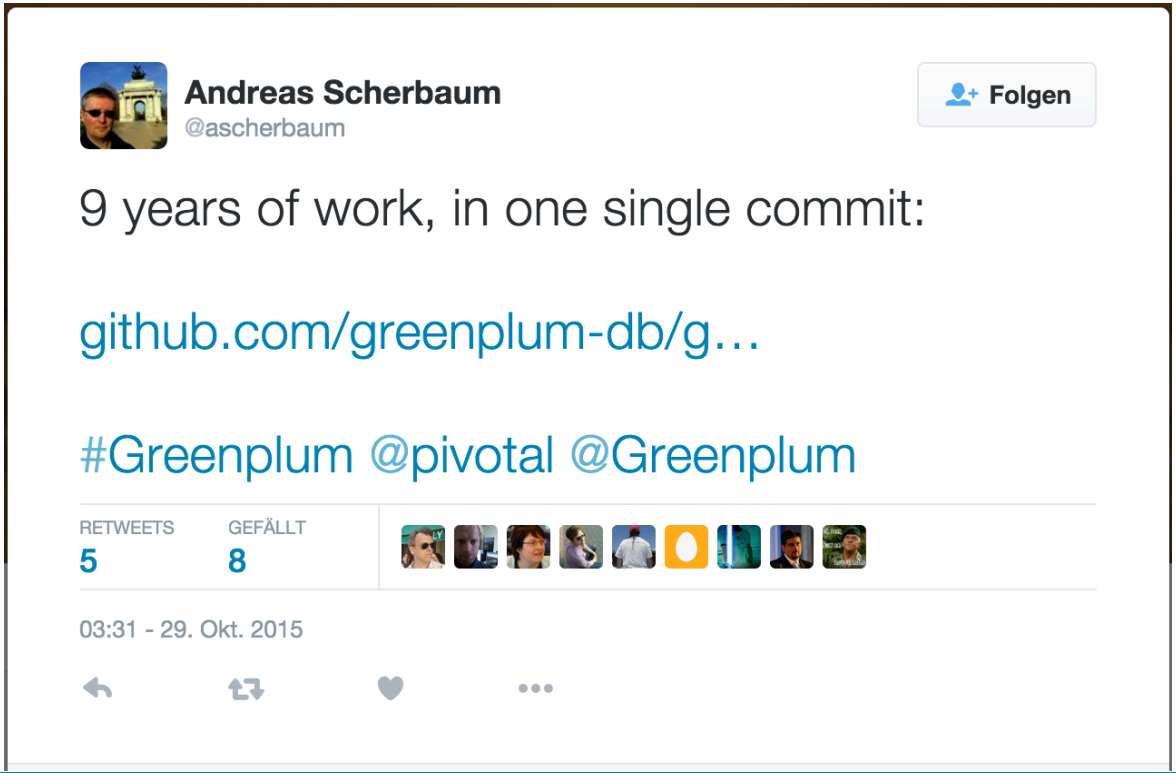

Pivotal

Greenplum Open Source

- Lives on GitHub
- Based on PostgreSQL 8.2 commits
  - Preserves the history

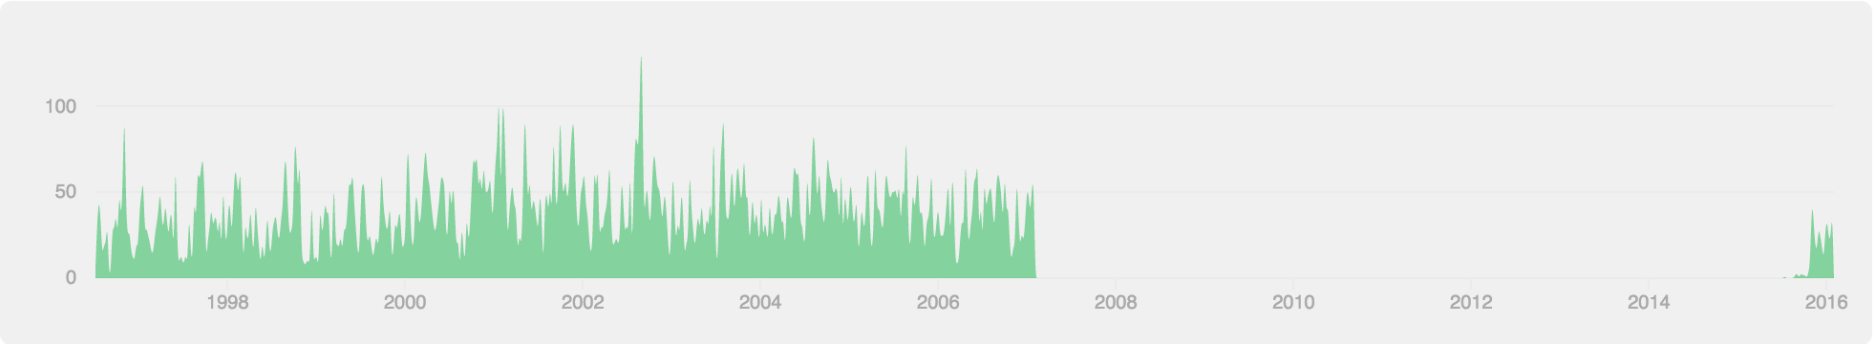

Pivotal

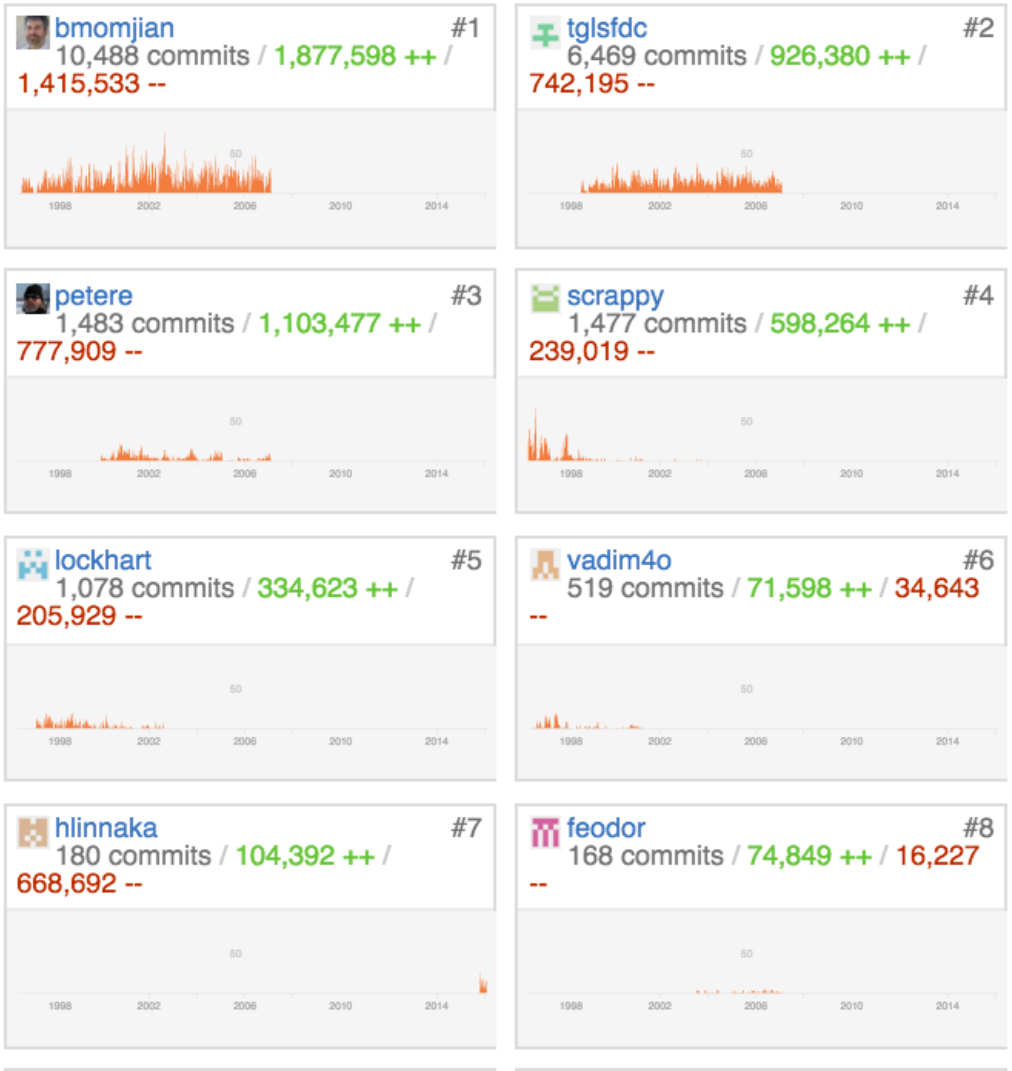

Pivotal

History of Greenplum - 2016

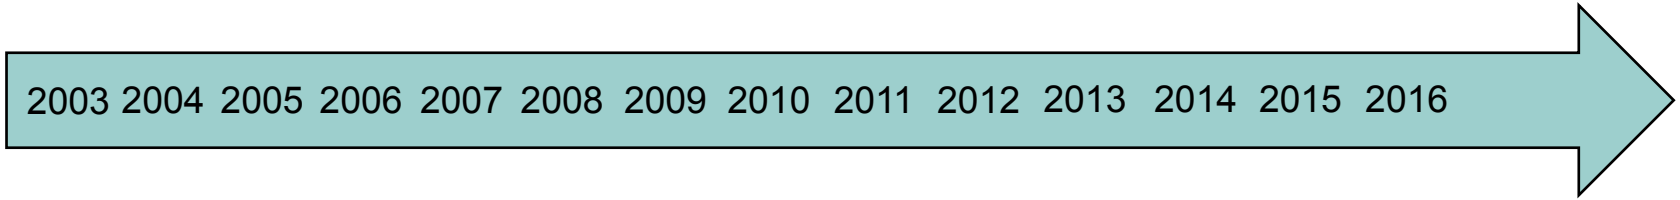

History of Greenplum - Future

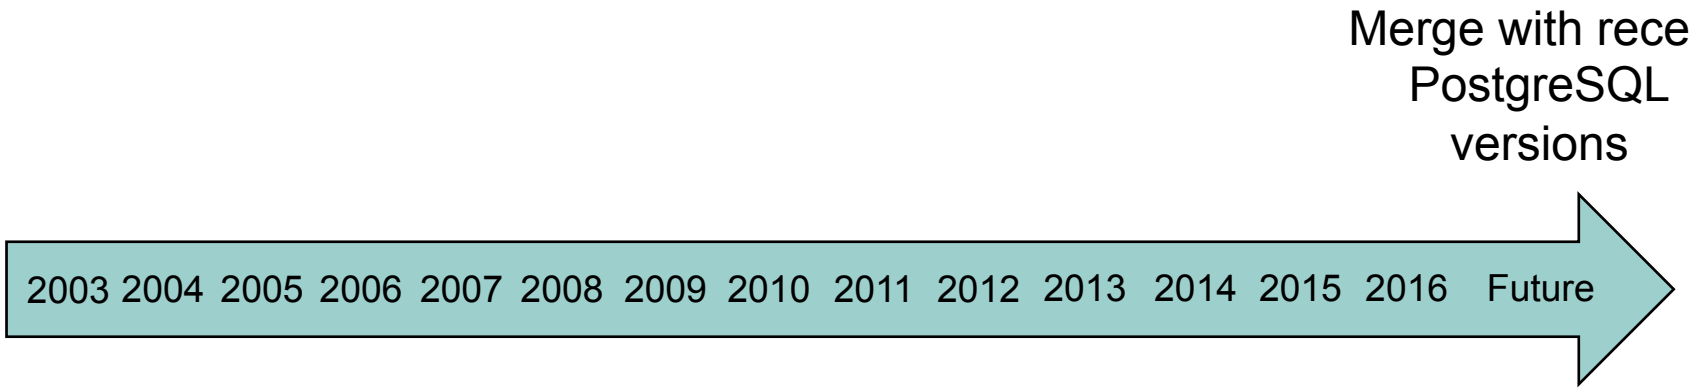

Why Open Source?

Why go Open Source?

- “Opening things up is incredibly healthy for the company”
  - For all products
  - Forced us to clean things up
  - Having to justify the decisions help in better communicating internally and externally

Pivotal

Why go Open Source?

- Customers more and more request Open Source products in RFPs
- Customers do not want to be locked in into one vendor or product (anymore)
- A healthy community improves the development process
- Restructure of the code to provide better development workflow

Pivotal

Why go Open Source?

- And of course: People live Open Source

Pivotal

Challenges

Roadblocks along the way …

Pivotal

Challenges

- How much time do you have?

Pivotal

Challenges

- License
  - Which license is compatible and sufficient
  - Settled with Apache License: PostgreSQL License has no patent protection
- Patents
  - Existing patents can’t be ignored
- Project owner / Governance model
  - Make it an ASF project?
  - Give it to the PostgreSQL Project? (they refused)

Pivotal

Challenges

- Customer names in code
  - Quite a challenge with ~1000 current customers
- Dependencies on internal build system
  - Only available in VPN and office network
- One single repository
  - Perforce, Copies in Atlassian Stash

Pivotal

Challenges

- Test systems
  - Test systems incomplete
  - Too many test systems
  - Too complex test systems
  - Test systems too big (tests run for several hours)
  - Most test systems only accessible for Pivotal

Challenges

- CVEs
  - Followed up with each and every PostgreSQL CVEs
- Static code scans
  - Two different products
- Black Duck scans
  - Scans for license violations

Pivotal

Challenges

- Caring parents …
  - EMC is not known for being an Open Source company
- Split time between going open source and handling customer escalations
- Open Source as a strategy, business model
- What is the IP portfolio?

Pivotal

Challenges

- Tension between commercial interests and open source interests
- PostgreSQL

Pivotal

Patent warning about the Greenplum source code

From: Bruce Momjian <bruce(at)momjian(dot)us>
To: PostgreSQL-development <pgsql-hackers(at)postgreSQL(dot)org>
Subject: Patent warning about the Greenplum source code
Date: 2015-10-30 08:47:35
Message-ID: 20151030084735.GA23947@momjian.us (view raw or whole thread)
Thread: 2015-10-30 08:47:35 from Bruce Momjian <bruce(at)momjian(dot)us>
Lists: pgsql-hackers

Some of you might have seen that the Greenplum database source code has
been published:

        https://adtmag.com/articles/2015/10/28/greenplum-open-sourced.aspx

under the Apache 2.0 license:

        http://www.apache.org/licenses/LICENSE-2.0

The source code has known patents owned by Pivotal/Greenplum. The
license has a patent grant clause:

Challenges

- Documentation is partly open, partly closed
  - Updating documentation is a challenge
- Culture change: employee -> community
  - Have every product related discussion in the public
- Have code open for inspection by customers, competitors
  - They can find problems, bugs, or just copy ideas

Pivotal

Lawyers

- You need them in the process. Yes, really.
- “Programmers all think they are lawyers”

Lessons learned

History repeats itself ...

Pivotal

Think twice before you fork

- Overwhelming opinion today is that forking PostgreSQL (and not merging new versions) was a mistake
  - Over time, more and more desired features showed up in PostgreSQL

Pivotal

Think early about where you go

- Choice between Apache Software Foundation and PostgreSQL Project
  - ASF is not a good fit
  - PG did not want to “adopt” GP
  - Quote: in PG everything is just opinion (“503 in Canada”)

Pivotal

Development workflow

- GitHub is cool, but workflow is a bit different
  - PostgreSQL (and Greenplum) do not merge
- We missed setting up the Contributor License Agreement (CLA) process
  - Explain the CLA in detail!
  - This is not about signing over code to Pivotal exclusively.

CLA not provided

gpdbdreddbot commented on Nov 5, 2015

Hello andreasscherbaum!

Thanks for submitting this pull request!

All pull request authors must have a Contributor License Agreement (CLA) on-file with us. Please sign the appropriate CLA (individual or corporate).

When sending signed CLA please provide your github username in case of individual CLA or the list of github usernames that can make pull requests on behalf of your organization.

If you are confident that you're covered under a Corporate CLA, please make sure you've publicized your membership in the appropriate Github Org, per these instructions.

Pivotal

CLA provided

gpdbdreddbot commented on Nov 11, 2015

Hello andreasscherbaum!

Thanks for submitting this pull request! I'm here to inform the recipients of the pull request that you've already signed the CLA.

Pivotal

The small things

- Domain(s)
  - We literally got greenplum.org back days before the release
- Website
  - You think it's still time – until it's too late
- Mailinglists
  - Which ones, and how many
  - Have everyone in your company subscribed to the lists
- Talk with your lawyers, frequently

Pivotal

Go public

- We decided early to announce at pgconf.eu 2015
  - Choice between pgconfSV and pgconf.eu
  - Gives you a fixed timeline – which is good and bad
- Have swag ;-)
  - Marketing supported the Open Source announcement with nice merchandising articles

Pivotal

Q & A

Time for your questions ...

Pivotal

Thanks to (in no particular order)

- Mike Waas
- Lyublena Antova
- John Eshleman
- Ivan Novick
- Roman Shaposhnik
- Elisabeth Hendrickson
- Caleb Welton
- Ed Espino
- George Tuma

Pivotal

Upcoming conferences

- PostgreSQL Conference Europe
  - Tallinn, Estonia
  - November 1 – 4, 2016
  - Radisson Blu Hotel Olümpia
  - http://2016.pgconf.eu/

Pivotal


Pivotal®