POSTGRES UNIVERSAL DATABASE

Oleg Bartunov
Moscow University
Postgres Professional

CERN, Jan 17, 2020

Postgres developer/contributor since 1995

When I started using Postgres

- No Slonik
- No UTF-8, even no 8-bit
- No WAL
- No MVCC
- No replication
- No usable non-scalar data types
- No subselects, no window functions, no CTE

- It was Postgres95

OLEG BARTUNOV

PROFESSIONAL Postgres

Postgres developer/contributor since 1995

OLEG BARTUNOV

PROFESSIONAL Postgres

How to choose a right database ?

- **People usually choose a database looking on**
  - Functionality, Performance
  - Availability - License, price
  - Local expertise, Personal experience
  - Compatibility to existing environment
  - Support
- **After project started**
  - Need new functionality, Better performance
- **Project is in production, no way to change database**
  - Starting to use various ugly «solutions»
  - System works, but looks pretty strange

If you chose a wrong database

System works, but looks pretty strange

PostgreSQL Universal Database

- Any project could start with PostgreSQL
- PostgreSQL is a reliable and stable database with rich functionality and long history
- PostgreSQL has liberal BSD license, cross platform (~30)
- Developed by international community, no vendor lock

- PostgreSQL is **EXTENSIBLE**, this is the very important feature, which people miss ! It allow database to support
  - New workloads
  - New functionality
  - New environment
  - Often without restarting a server, no need core programmer.

PostgreSQL Universal Database

Extensibility makes PostgreSQL Universal Database !

“It is imperative that a user be able to construct new access methods to provide efficient access to instances of nontraditional base types”

Michael Stonebraker, Jeff Anton, Michael Hirohama.

Extendability in POSTGRES , IEEE Data Eng. Bull. 10 (2) pp.16-23, 1987

Postgres can be extended:

Functions, data types, operators

Procedural languages (sql, pl/pgsql, pl/perl, pl/python, pl/tcl, pl/R, pl/java, …,pl/v8)

Indexes (Btree, Hash, GiST, GIN, SP-GiST, BRIN, BLOOM)

Foreign Data Wrappers (almost to all databases)

PostgreSQL: OLTP, MPP, OLAP, CLOUD, GIS, STREAM, TIMESERIES, GPU, NoSQL

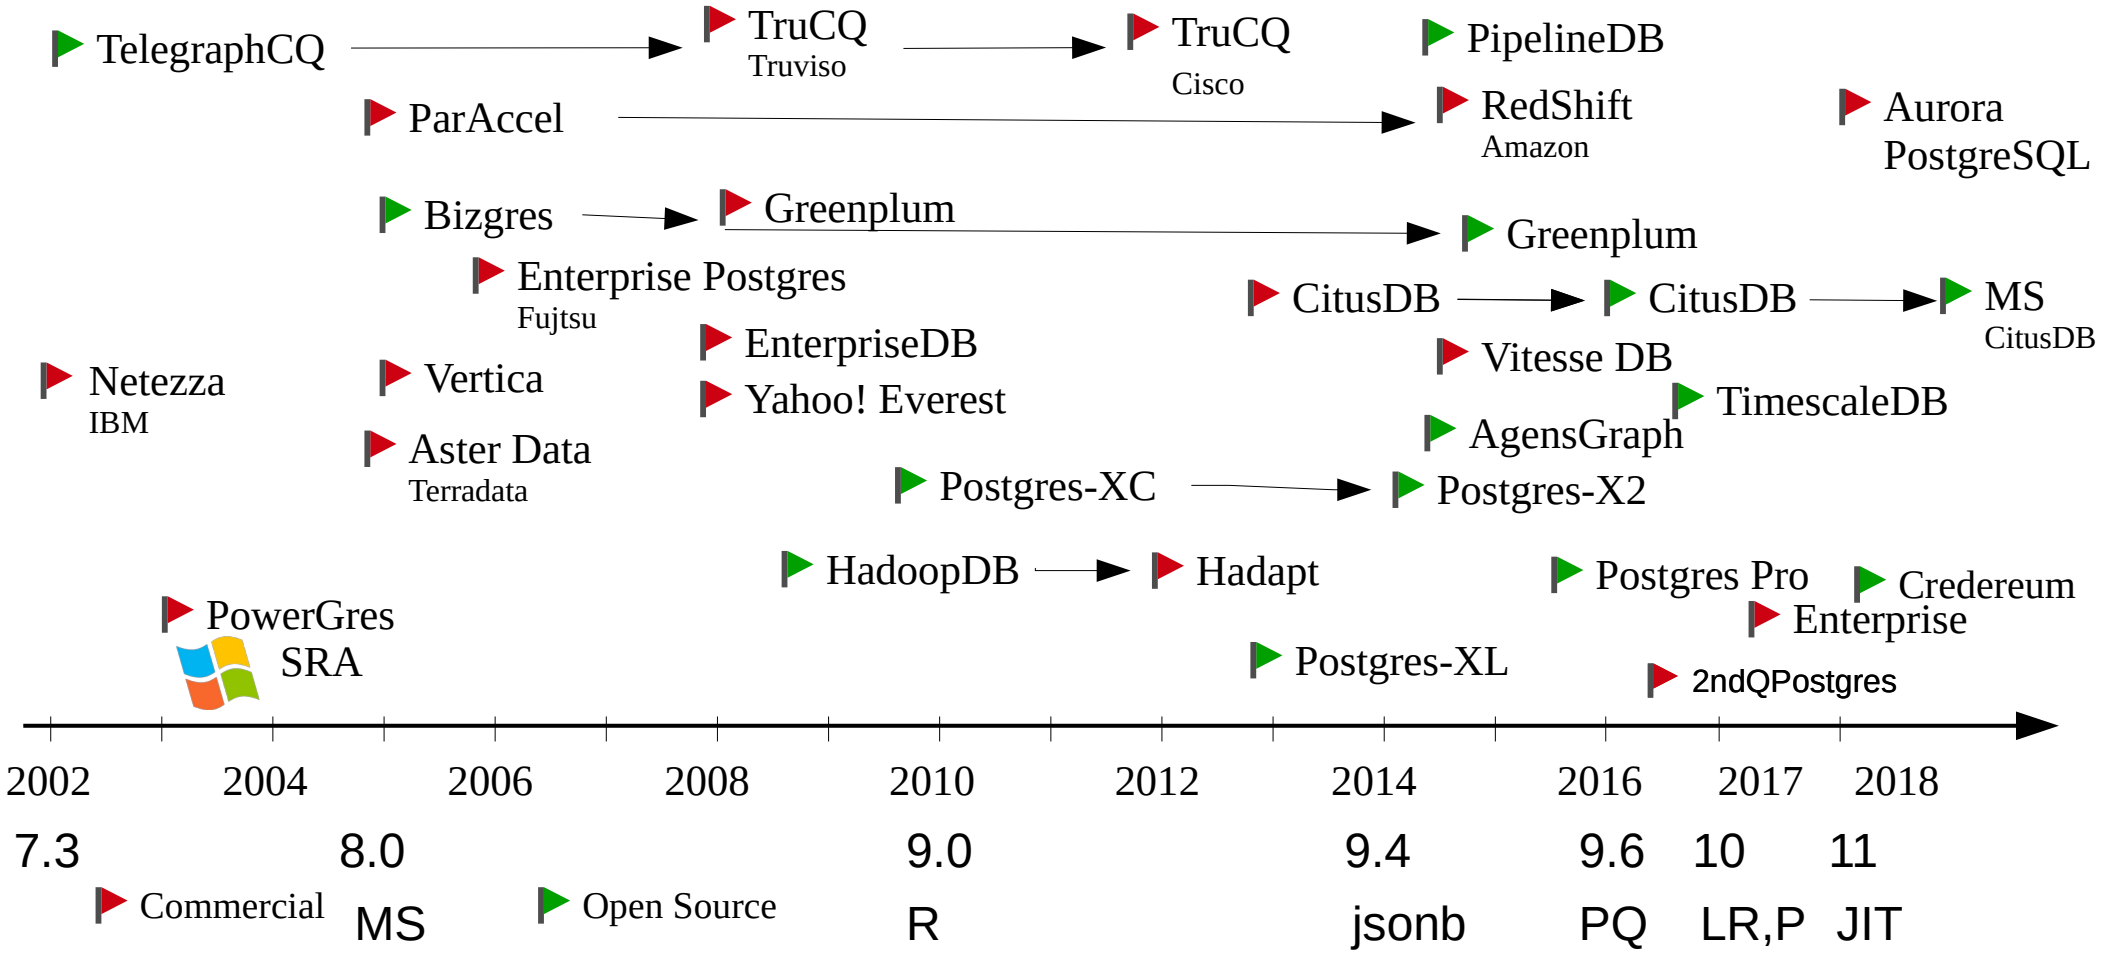

Postgres Evolution

1996 — Project starts
1997 (6.1) — Internationalization **+World**
2005 (8) — Windows support +**Window users**
2010 (9) — Built-in replication +**Enterprise users**
2014 (9.4) — Jsonb +**NoSQL users**
2016 (9.6) — Parallel Query +**OLAP users**
2017 (10) — Logical Replication, Declarative Partitioning
2018 (11) — JIT
2019 (12) — Pluggable storage API, SQL/JSON
202X (?) - Cloud support ! +**ALL**

Postgres Community Evolution

- 198X - Academic Postgres (x10)
- 1995 - Community Postgres95 (<400)
- PostgreSQL V6

  Community develops for Community
- 200X — First Postgres-centric companies

  (GreatBridge, 2ndQuadrant, EDB...)

  +Full-time developers for Community

  - First enterprise forks

Professional Postgres

- 2010 — Enterprise companies recognize Postgres
- 2015 — Majority of major developers were hired by PG-companies (+Citus Data, +Postgres Professional)
  - Now the companies drive the development
  - Community: test, approve
  - Postgres became Enterprise ready (More forks)
  - Postgres became Professional

PG-companies drive the development

PG-companies - proxy between Enterprise and Community

- Big enterprises require additional features "right now"
- PG-companies develop, support and test these features in their forks
- Some features returned back to community (not easy)
- Community accept (if) and support code

Example: Postgres Pro Enterprise

- 64-bit XID (community?)
- Multi Master cluster (community?)
- Incremental backup (pg_probackup, opensource)
- Advanced partitioning (pg_pathman, opensource)
- Threaded Postgres (community?)
- Sharding (in development)
- Built-in pooler (PG13 ?)
- Seamless upgrade (in development)
- SQL/JSON (PG12)

Popularity of PostgreSQL is growing

DB-Engine

Hacker News Hiring Trends - 2019

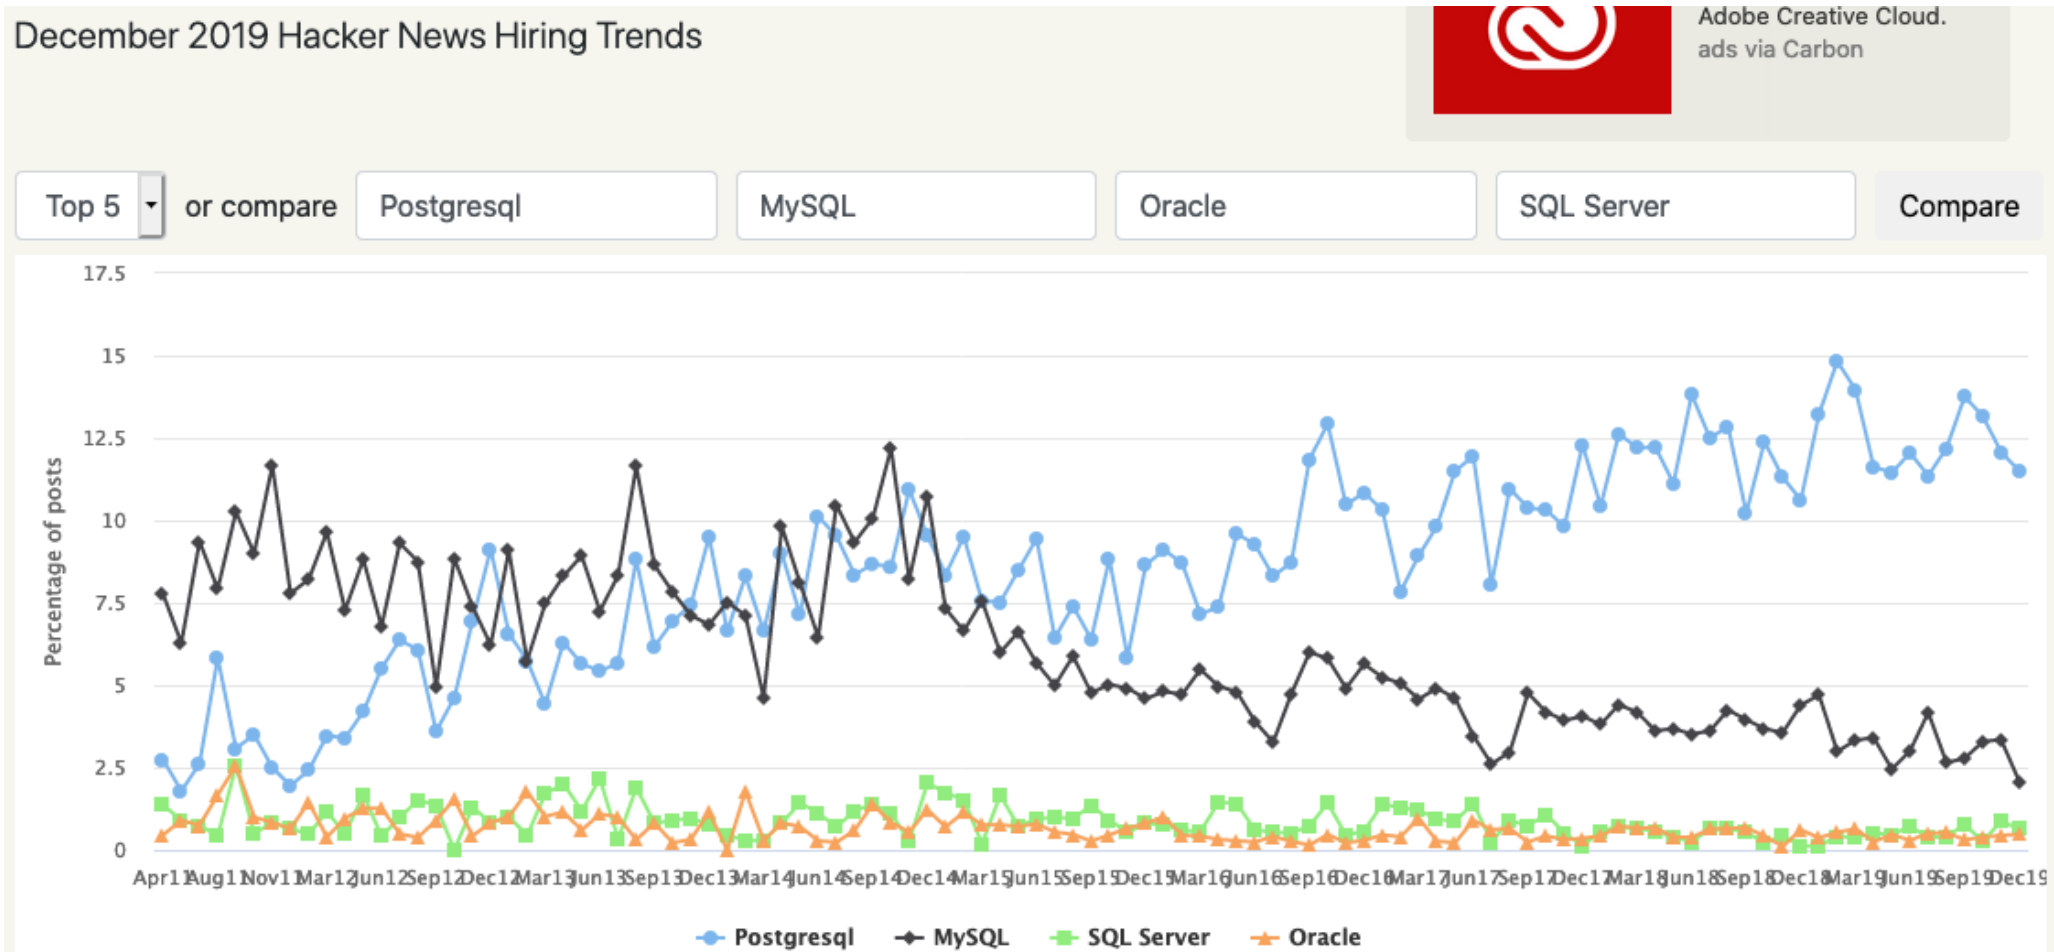

Postgres PROFESSIONAL Stack Overflow developer survey 2019

MOST USED

MOST LOVED

Databases

All Respondents | Professional Developers

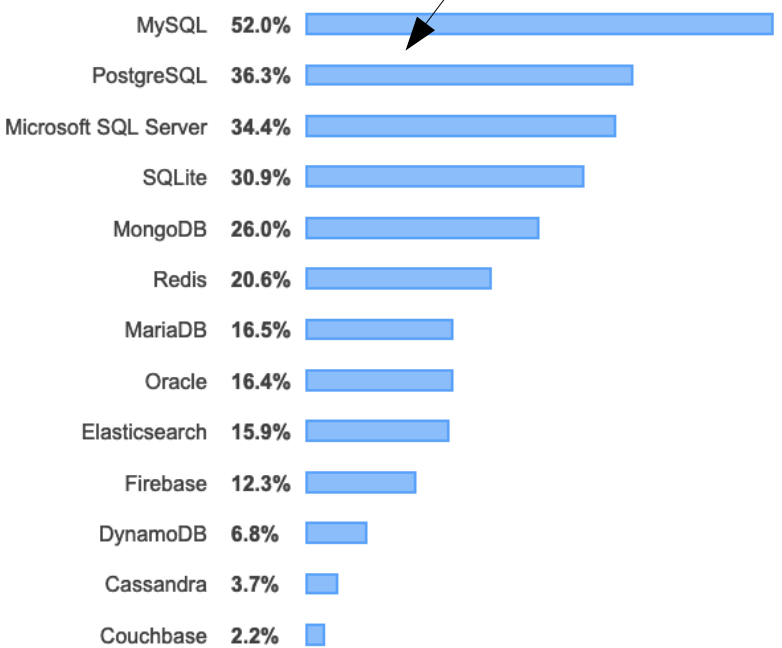

Most Loved, Dreaded, and Wanted Databases

Loved | Dreaded | Wanted

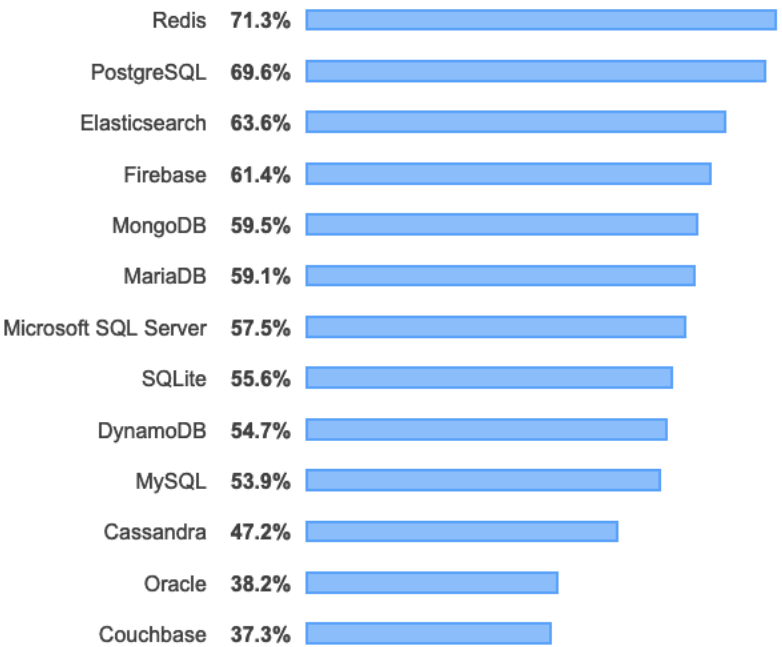

Postgres in Russia is database #1

Postgres PROFESSIONAL

PostgreSQL > Oracle in Russia

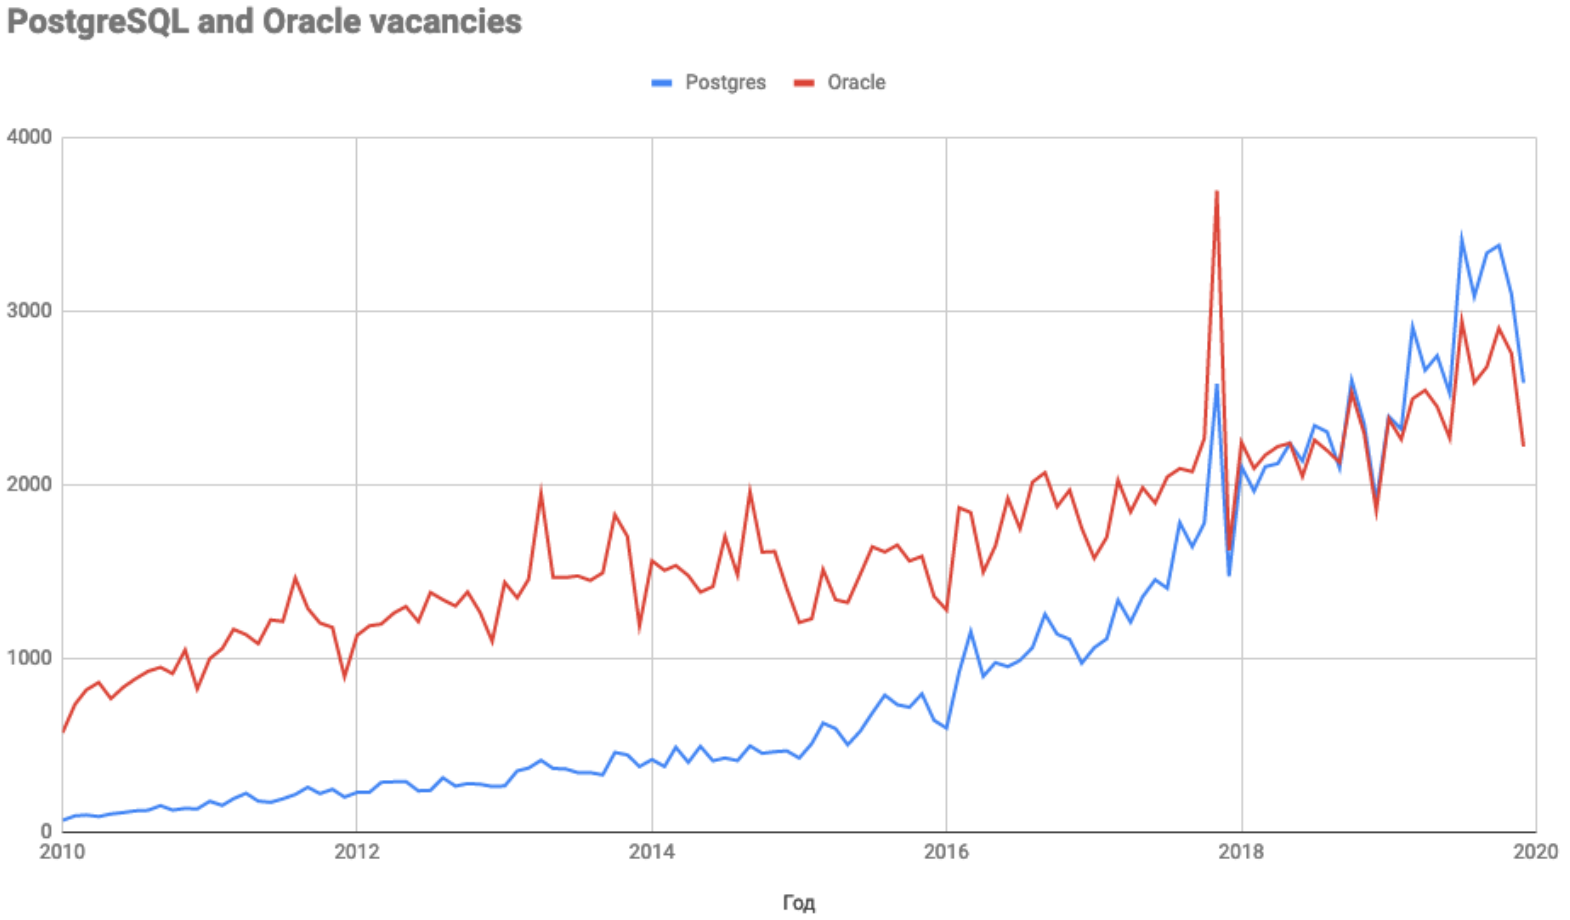

Web and Postgres: 1996-2018

- 1996: Start using Postgres on Web, no 8-bit support — introduced **locale** support
- 1999: World's top-5 portal. We start with PostgreSQL 6.5. on server ~ my smartphone to support > 1 mln. users/day. Quickly run out of resources
- Denormalize, use arrays -> slow -> discover **GiST** → improve **GiST** - **intarray** with indexes
- Need **FTS**, made **tsearch** using intarray and GiST indexes
- Need fast search on hierarchical data — **ltree** — GiST indexes
- Need flexible schema — **hstore** — GiST index
- Need faster FTS — **GIN** index for tsearch, hstore
- Need misprint search — **pg_trgm** — GiST/GIN indexes
- Compete to NoSQL - better/binary json - **jsonb** — GIN index - **jsonpath**
- Need faster FTS — **RUM** access methods

NoSQL Postgres

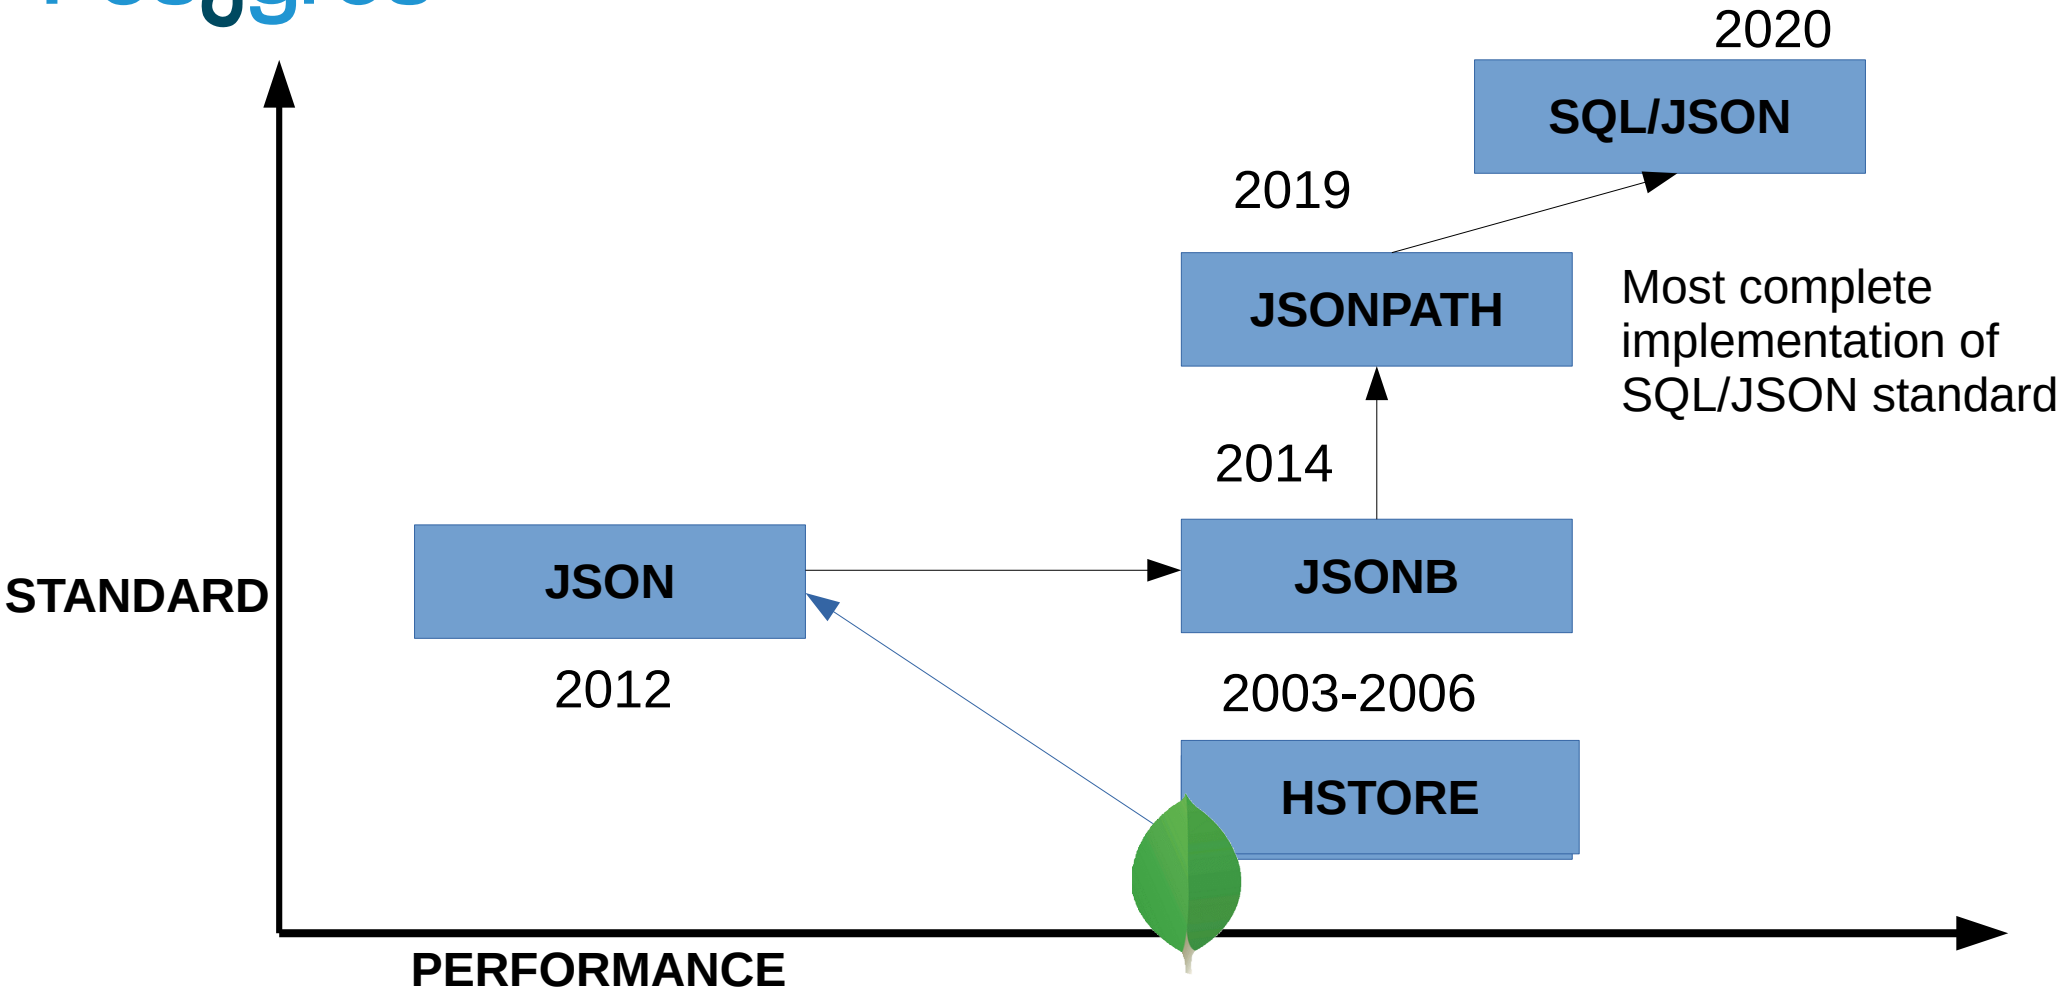

PROFESSIONAL Postgres

What is the fate of the Universe ?

WRITTEN IN THE STARS

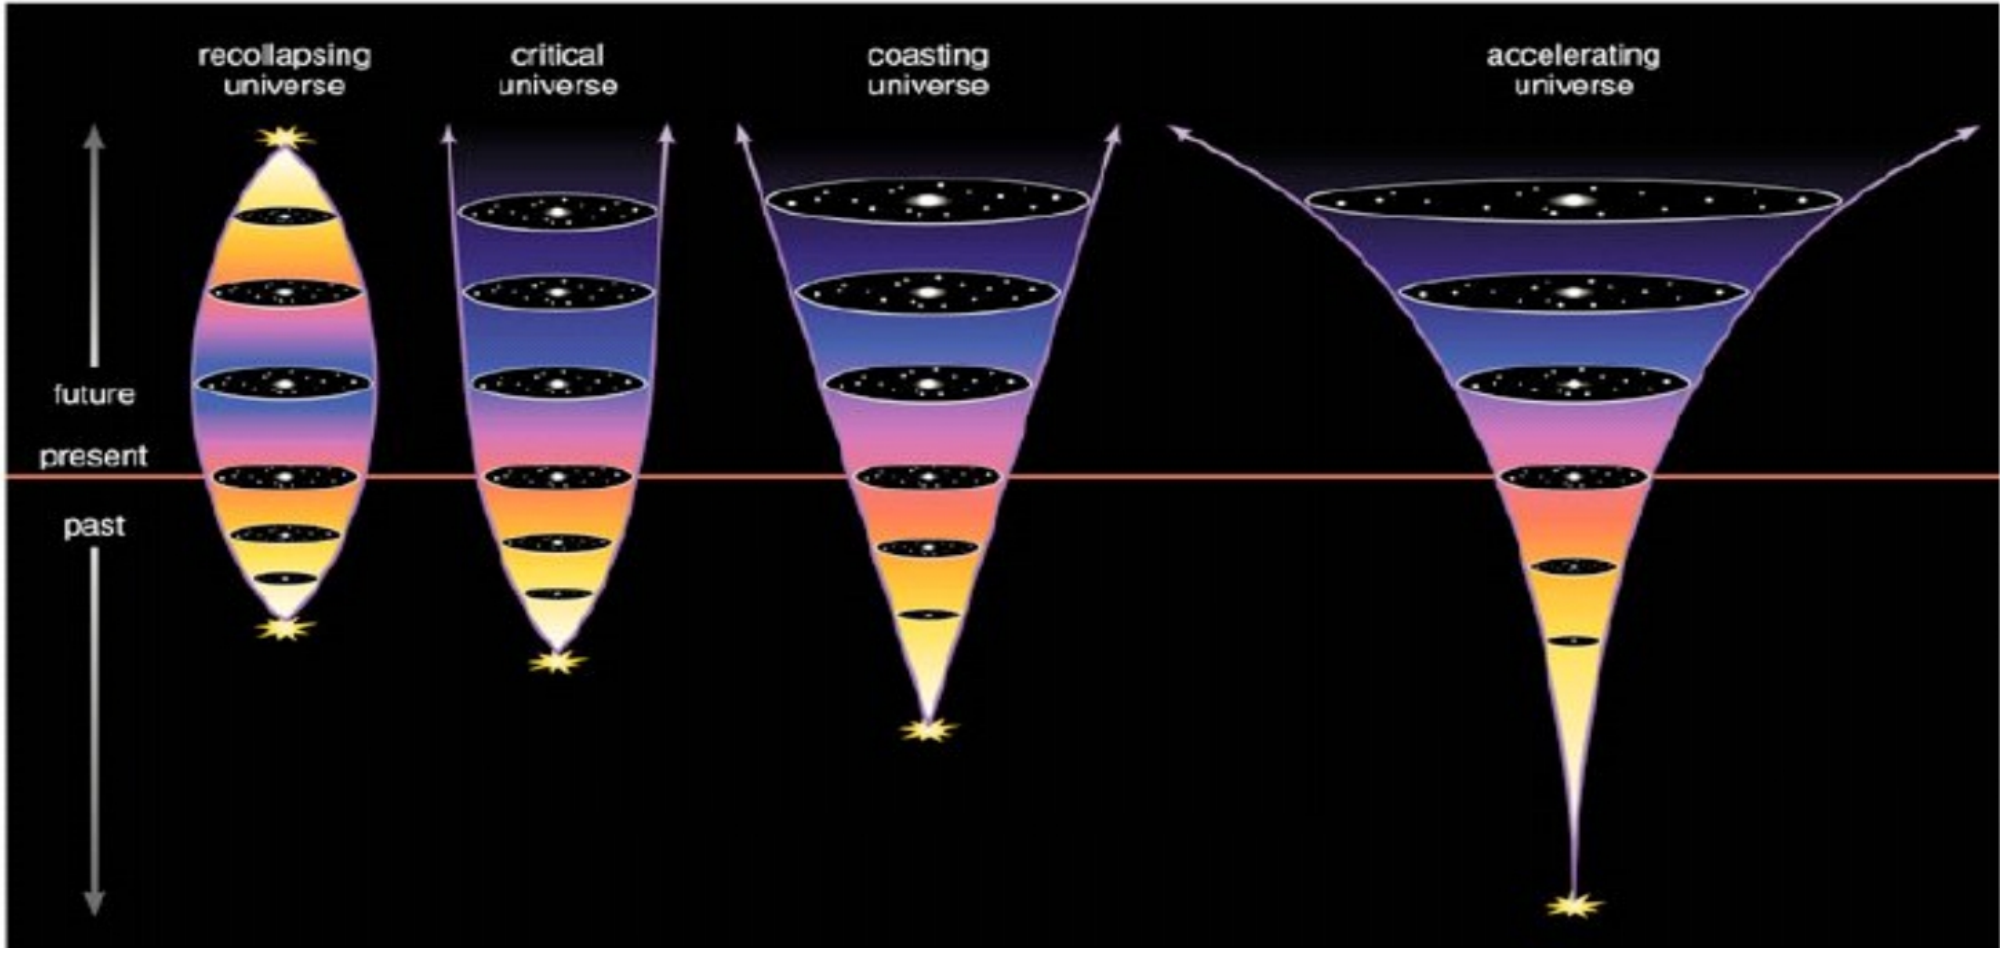

Postgres PROFESSIONAL

SN Ia 1994D in NGC 4526

PROFESSIONAL Postgres

Supernovae 1987A

SN 1987a, Type II, +2.9, Tarantula Nebulae in LMC, 168 000 ly, progenitor: Sanduleak -69° 202,

The Scale of the Universe

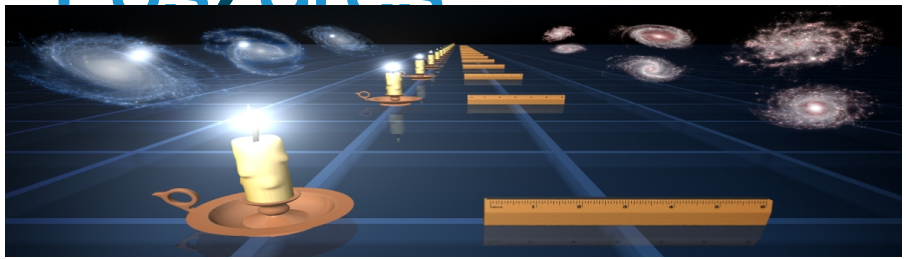

- Supernovae(Ia) - «standard candles»
- Used to measure the distance to the host galaxy

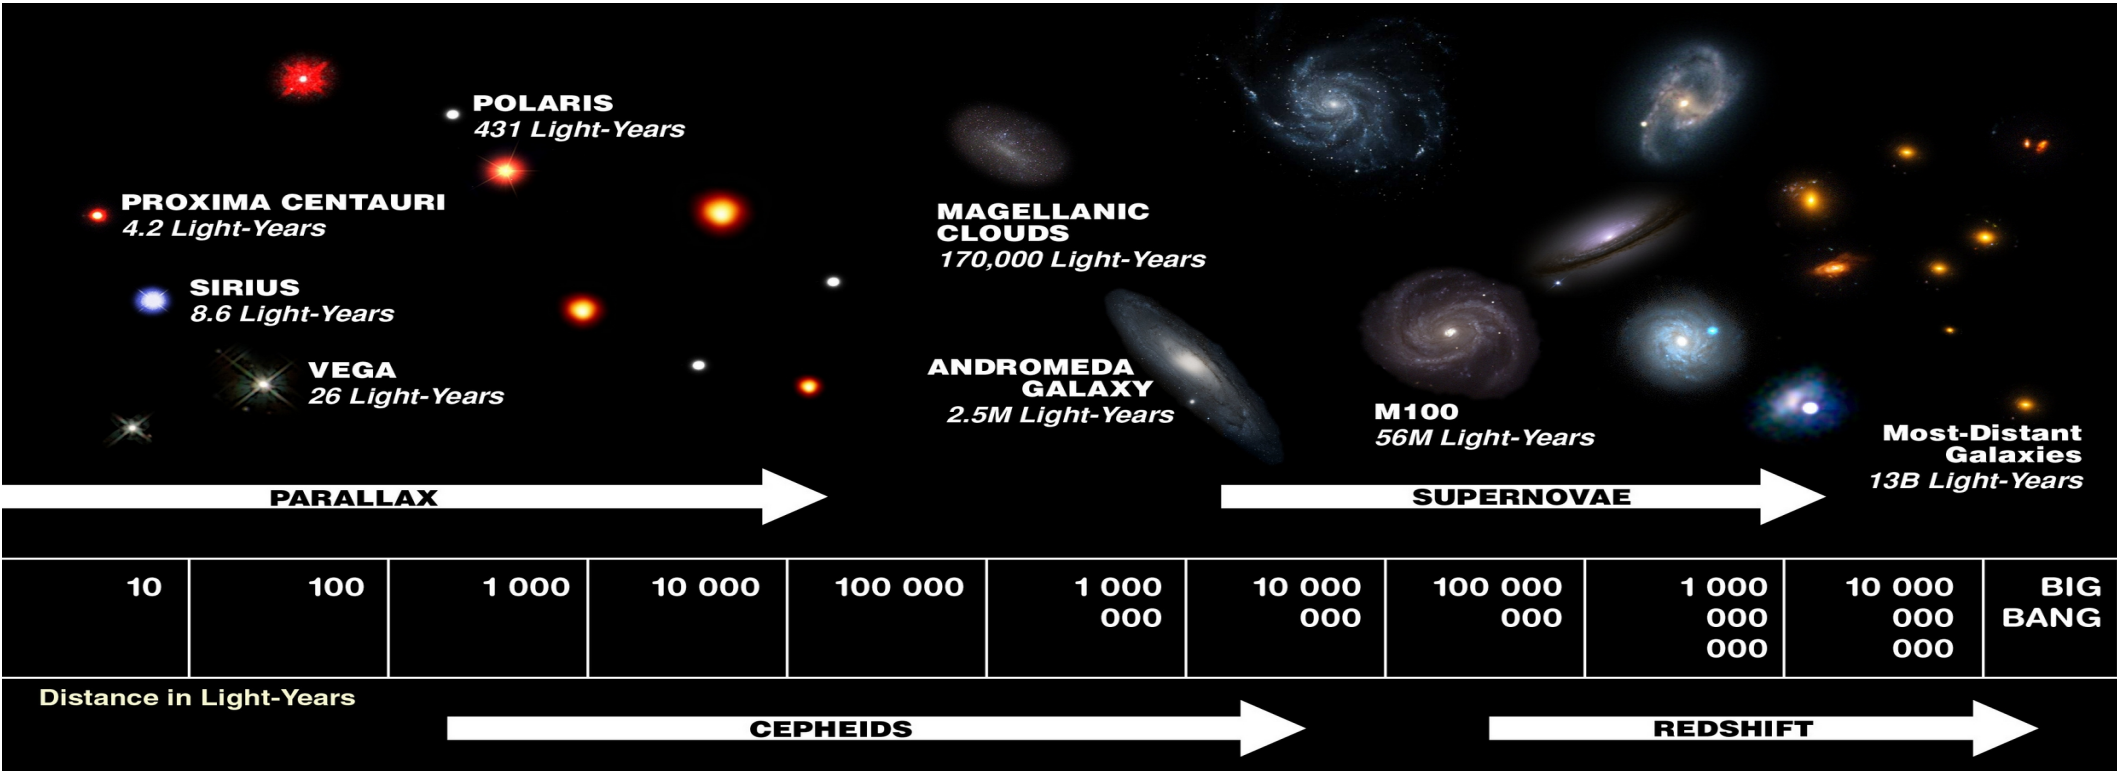

M31 (Andromeda), AZT-5

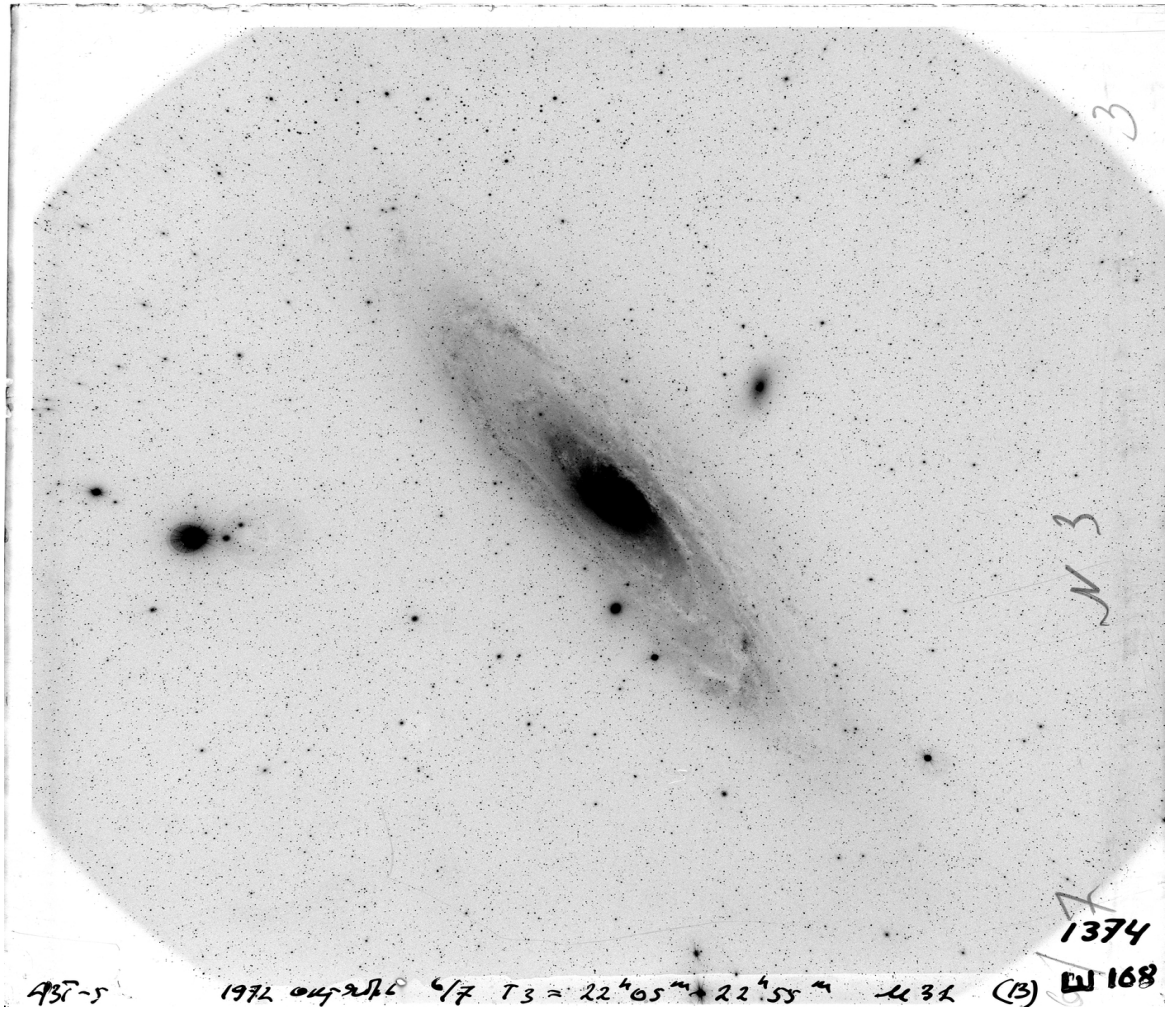

SN 2008fv in NGC 3147, Draco
Dmitry Tsvetkov, SAI MSU

PROFESSIONAL Postgres

Blink Comparator (Manual Discovery)

Many hours of hard work !

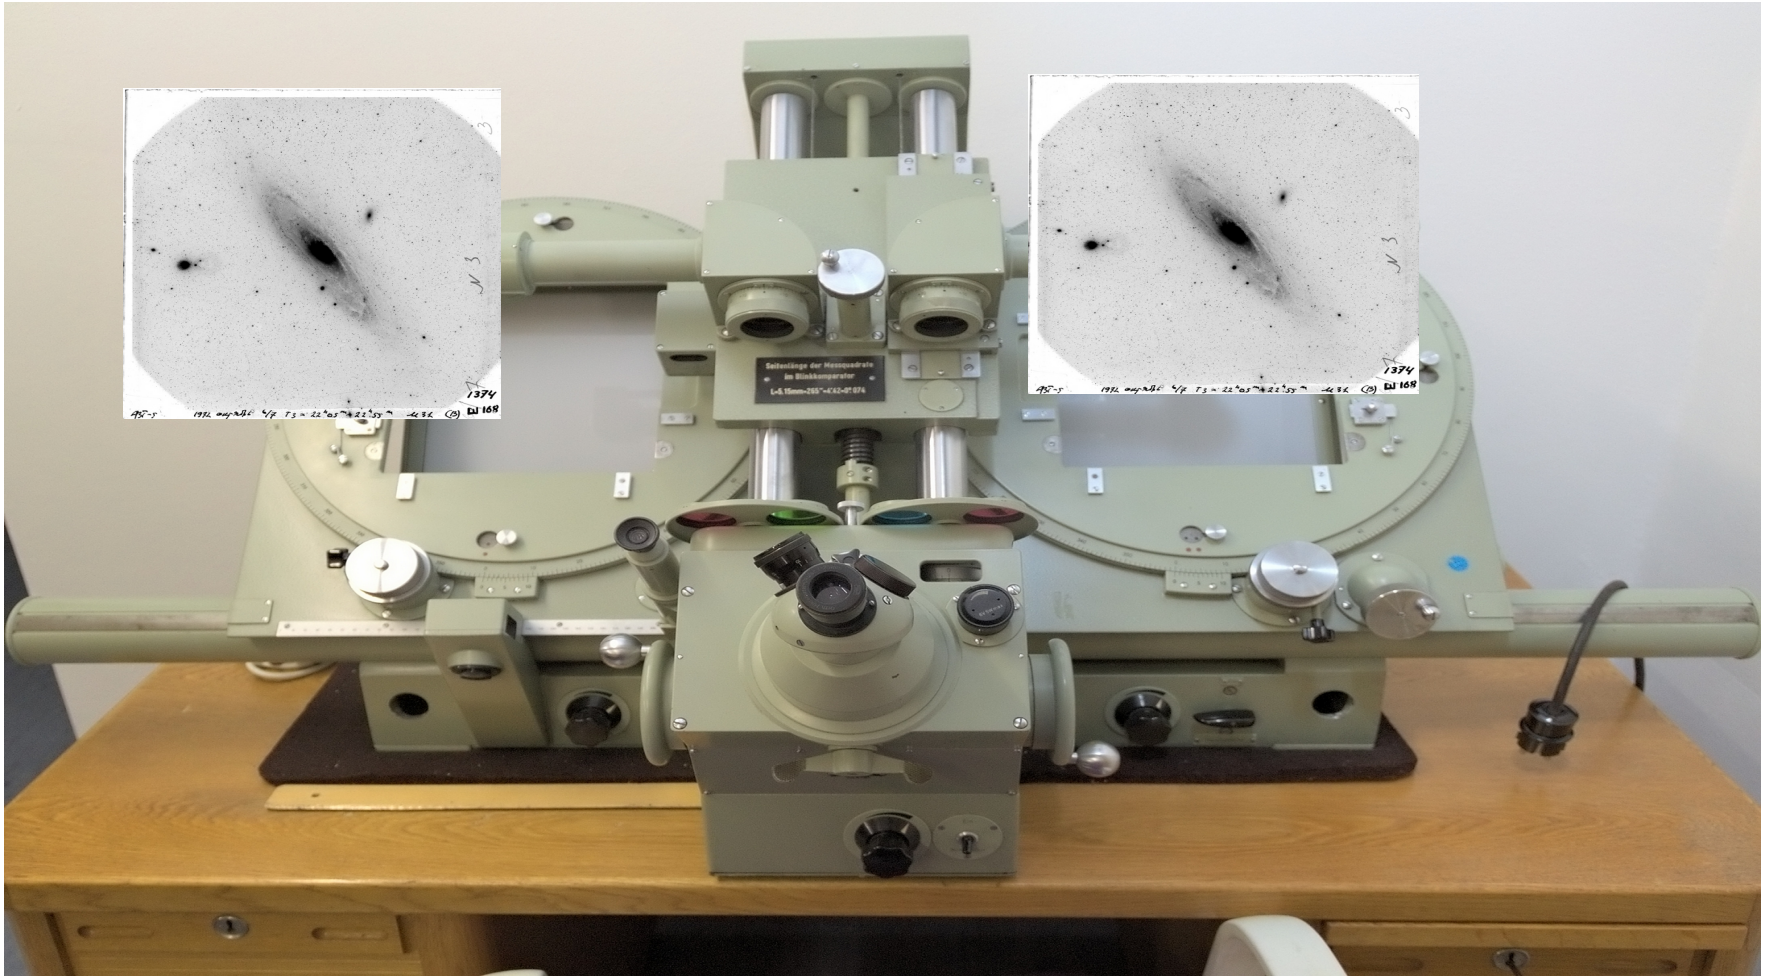

M31 (Andromeda), AZT-5

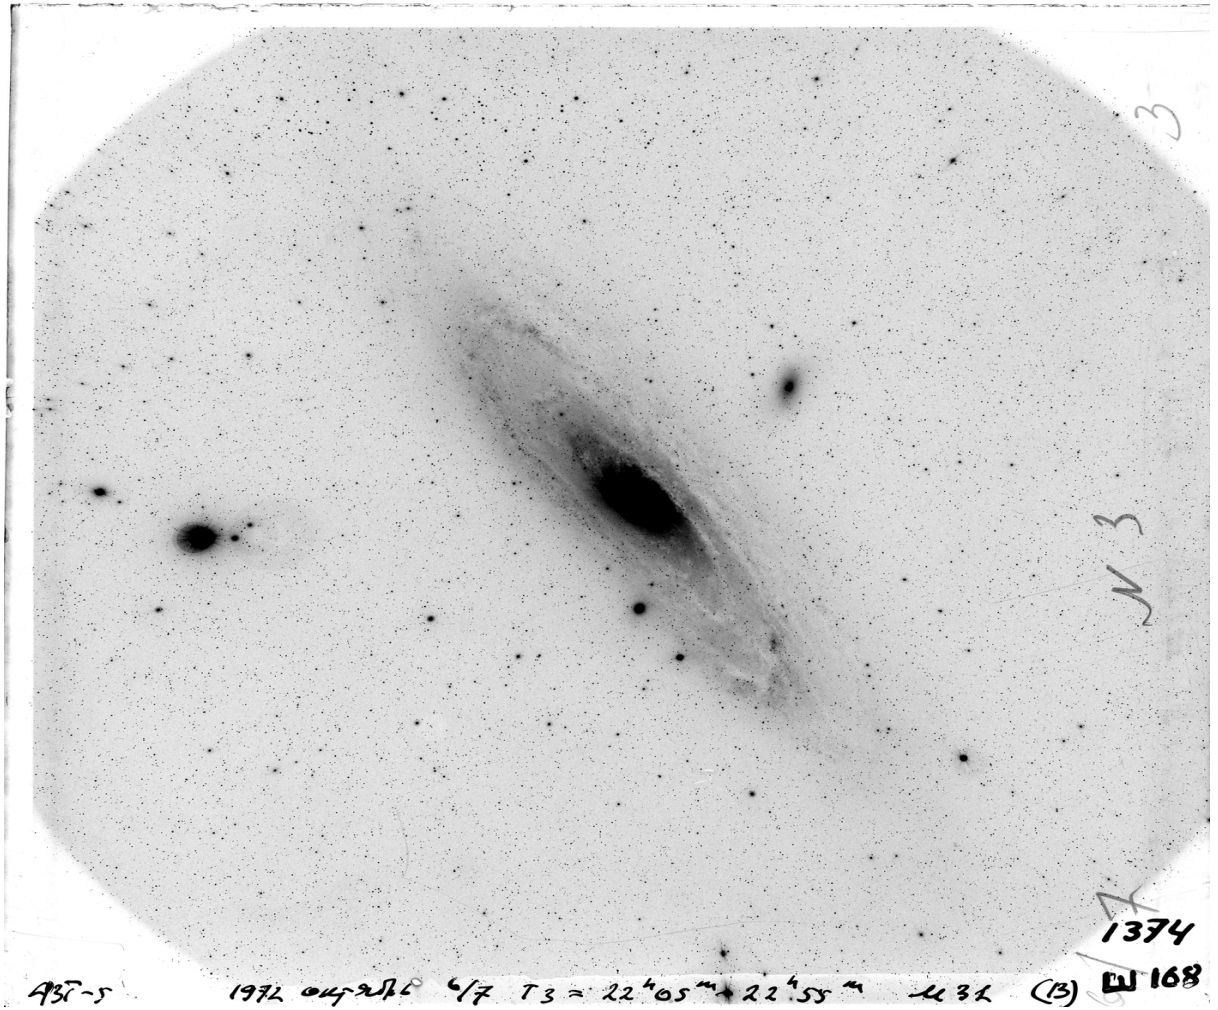


```
x-00001|1.1458447|-89.9186147
x-00002|1.3300139|-89.9332336
x-00003|3.2556022|-89.9641031
x-00004|3.6464625|-89.9060142
x-00005|6.3110253|-89.9523947
x-00006|6.6275517|-89.9279197
x-00007|7.8266025|-89.9129272
x-00008|9.0694378|-89.9714031
x-00009|9.6627953|-89.9244314
x-00010|10.0494292|-89.9705058
x-00011|10.4863922|-89.9699058
x-00012|11.0953692|-89.9016031
x-00013|11.3240233|-89.9344336
x-00014|11.7906064|-89.9070308
x-00015|12.0416581|-89.9300586
x-00016|12.0522308|-89.9002281
x-00017|12.2808536|-89.9107669
x-00018|13.0316142|-89.9214558
x-00019|13.8727033|-89.9577031
x-00020|14.6546639|-89.9191919
x-00021|18.3035981|-89.9447475
x-00022|18.5185631|-89.9446836
x-00023|19.8675597|-89.9836308
x-00024|20.9699533|-89.9226864
x-00025|21.6777744|-89.9256808
x-00026|23.3660669|-89.9036558
x-00027|24.2841308|-89.9516475
x-00028|24.3273161|-89.9202392
x-00029|24.5540458|-89.9246003
x-00030|24.5655172|-89.9122336
x-00031|26.3487519|-89.9460336
x-00032|26.5268008|-89.9311503
x-00033|26.6070808|-89.9271808
x-00034|27.4104919|-89.9768558
x-00035|27.8290442|-89.9304622
x-00036|28.5552036|-89.9199117
x-00037|29.4407347|-89.9762836
x-00038|30.5729608|-89.9377753
x-00039|30.7101131|-89.9105642
x-00040|33.2918250|-89.9106614
x-00041|33.4843678|-89.9442058
```

Spatial Join (Machine Discovery, < 1s)

Observations: 10^5

x-00001	1.1458447	-89.9186147
x-00002	1.3300139	-89.9332336
x-00003	3.2556022	-89.9641031
x-00004	3.6464625	-89.9060142
x-00005	6.3110253	-89.9523947
x-00006	6.6275517	-89.9279197
x-00007	7.8266025	-89.9129272
x-00008	9.0694378	-89.9714031
x-00009	9.6627953	-89.9244314
x-00010	10.0494292	-89.9705058
x-00011	10.4863922	-89.9699058
x-00012	11.0953692	-89.9016031
x-00013	11.3240233	-89.9344336
x-00014	11.7906064	-89.9070308
x-00015	12.0416581	-89.9300586
x-00016	12.0522308	-89.9002281
x-00017	12.2808536	-89.9107669
x-00018	13.0316142	-89.9214558
x-00019	13.8727033	-89.9577031
x-00020	14.6546639	-89.9191919
x-00021	18.3035981	-89.9447475
x-00022	18.5185631	-89.9446836
x-00023	19.8675597	-89.9836308
x-00024	20.9699533	-89.9226864
x-00025	21.6777744	-89.9256808
x-00026	23.3660669	-89.9036558
x-00027	24.2841308	-89.9516475
x-00028	24.3273161	-89.9202392
x-00029	24.5540458	-89.9246003
x-00030	24.5655172	-89.9122336
x-00031	26.3487519	-89.9460336
x-00032	26.5268008	-89.9311503
x-00033	26.6070808	-89.9271808
x-00034	27.4104919	-89.9768558
x-00035	27.8290442	-89.9304622
x-00036	28.5552036	-89.9199117
x-00037	29.4407347	-89.9762836
x-00038	30.5729608	-89.9377753

$\bowtie\theta$

Catalog(s): 10^9

t-0000001	1.1458447	-89.9186147	0.015	0.028
t-0000002	1.3300139	-89.9332336	0.050	0.110
t-0000003	3.2556022	-89.9641031	0.050	0.050
t-0000004	3.6464625	-89.9060142	0.204	0.224
t-0000005	6.3110253	-89.9523947	0.114	0.050
t-0000006	6.6275517	-89.9279197	0.098	0.150
t-0000007	7.8266025	-89.9129272	0.025	0.021
t-0000008	9.0694378	-89.9714031	0.200	0.200
t-0000009	9.6627953	-89.9244314	0.000	0.000
t-0000010	10.0494292	-89.9705058	0.050	0.228
t-0000011	10.4863922	-89.9699058	0.200	0.200
t-0000012	11.0953692	-89.9016031	0.050	0.259
t-0000013	11.3240233	-89.9344336	0.050	0.050
t-0000014	11.7906064	-89.9070308	0.159	0.131
t-0000015	12.0416581	-89.9300586	0.216	0.050
t-0000016	12.0522308	-89.9002281	0.050	0.050
t-0000017	12.2808536	-89.9107669	0.050	0.050
t-0000018	13.0316142	-89.9214558	0.152	0.120
t-0000019	13.8727033	-89.9577031	0.050	0.121
t-0000020	14.6546639	-89.9191919	0.050	0.069
t-0000021	18.3035981	-89.9447475	0.139	0.440
t-0000022	18.5185631	-89.9446836	0.057	0.268
t-0000023	19.8675597	-89.9836308	0.050	0.120
t-0000024	20.9699533	-89.9226864	0.050	0.050
t-0000025	21.6777744	-89.9256808	0.055	0.105
t-0000026	23.3660669	-89.9036558	0.050	0.135
t-0000027	24.2841308	-89.9516475	0.213	0.050
t-0000028	24.3273161	-89.9202392	0.550	0.999
t-0000029	24.5540458	-89.9246003	0.160	0.086
t-0000030	24.5655172	-89.9122336	0.205	0.050
t-0000031	26.3487519	-89.9460336	0.050	0.095
t-0000032	26.5268008	-89.9311503	0.335	0.245
t-0000033	26.6070808	-89.9271808	0.050	0.075
t-0000034	27.4104919	-89.9768558	0.094	0.090
t-0000035	27.8290442	-89.9304622	0.017	0.019
t-0000036	28.5552036	-89.9199117	0.050	0.115
t-0000037	29.4407347	-89.9762836	0.635	0.265
t-0000038	30.5729608	-89.9377753	0.314	0.170

Astronomy meets database

Indexing the SKY with PostgreSQL

HEALPIX

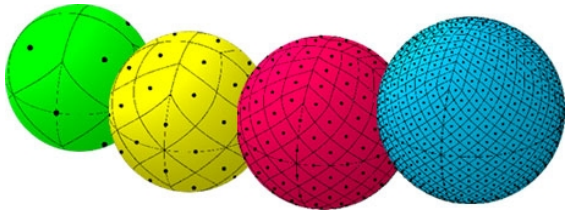

HTM(MS)

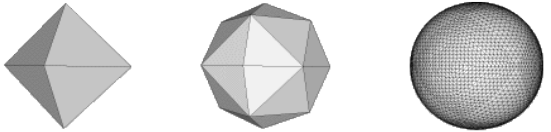

Q3C (PG)

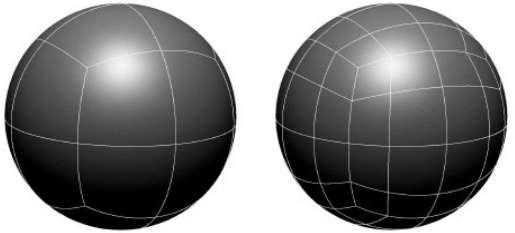

The sphere segmentation in Q3C

PGSphere (PG)

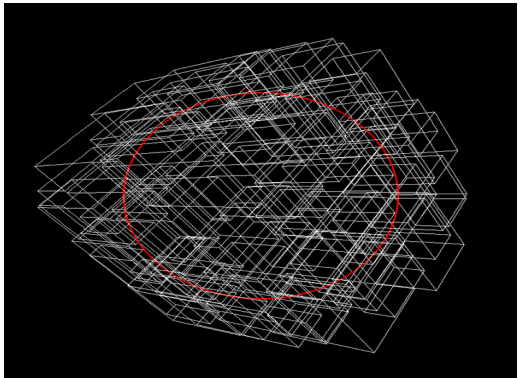

PROFESSIONAL Postgres

Acceleration Expansion !

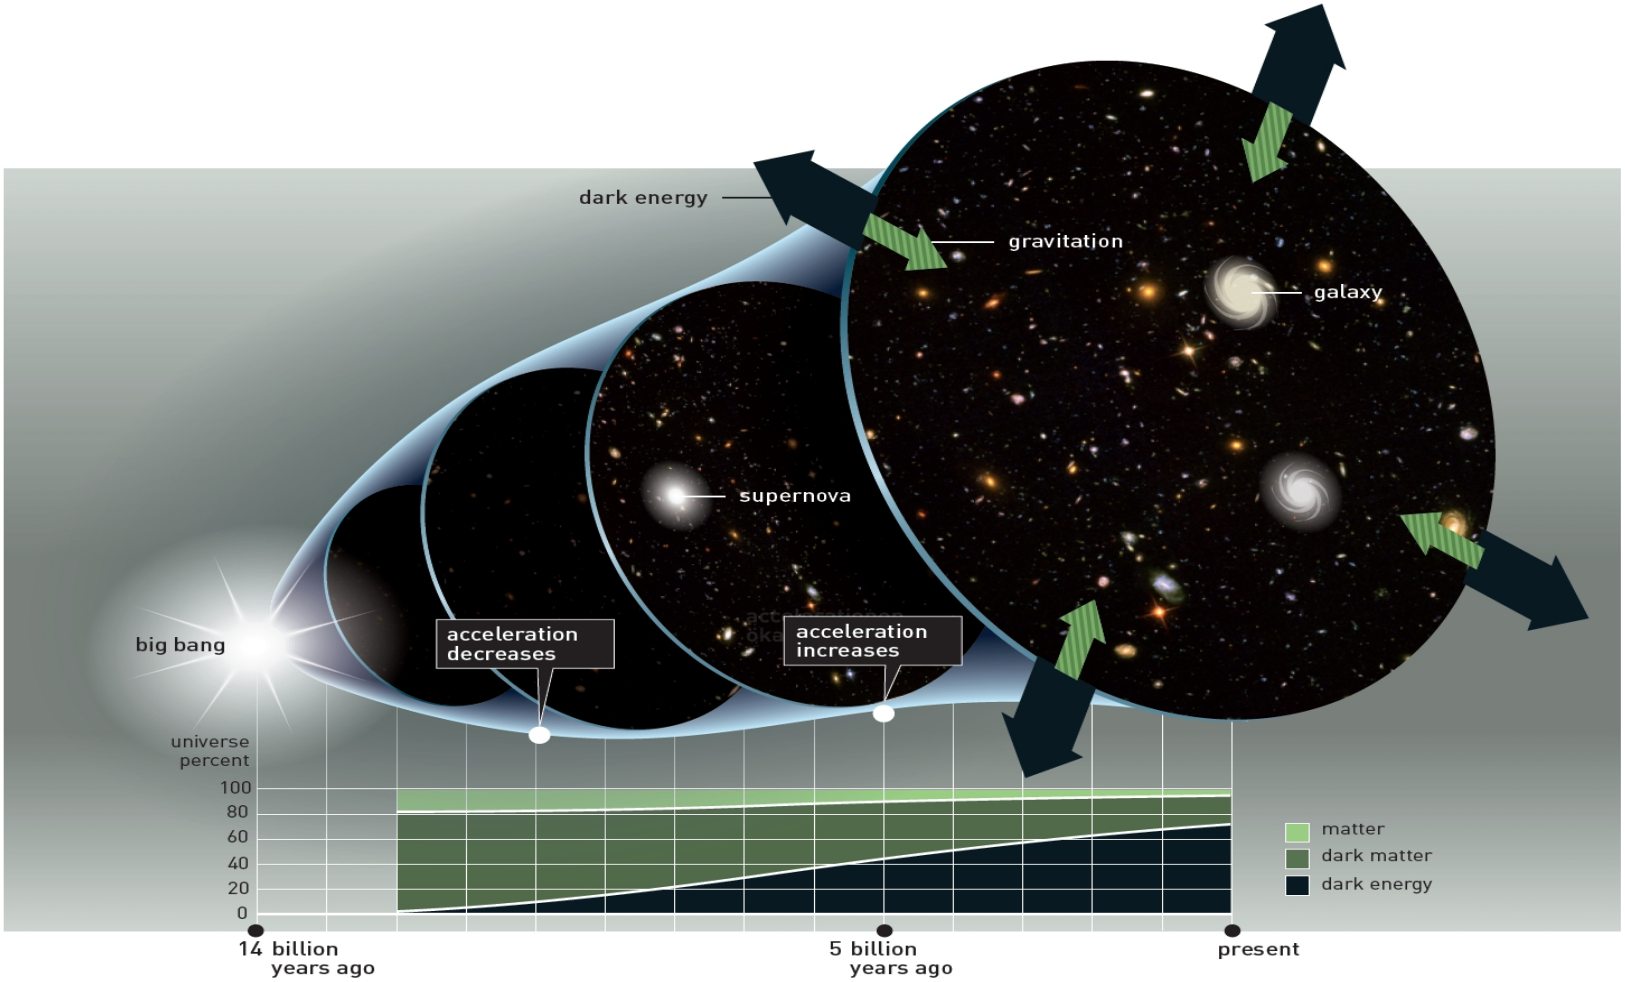

PROFESSIONAL Postgres

Nobel Prize in Physics 2011

Saul Perlmutter ,Brian P. Schmidt ,Adam G. Riess

Postgres in the Cloud

The future of Postgres

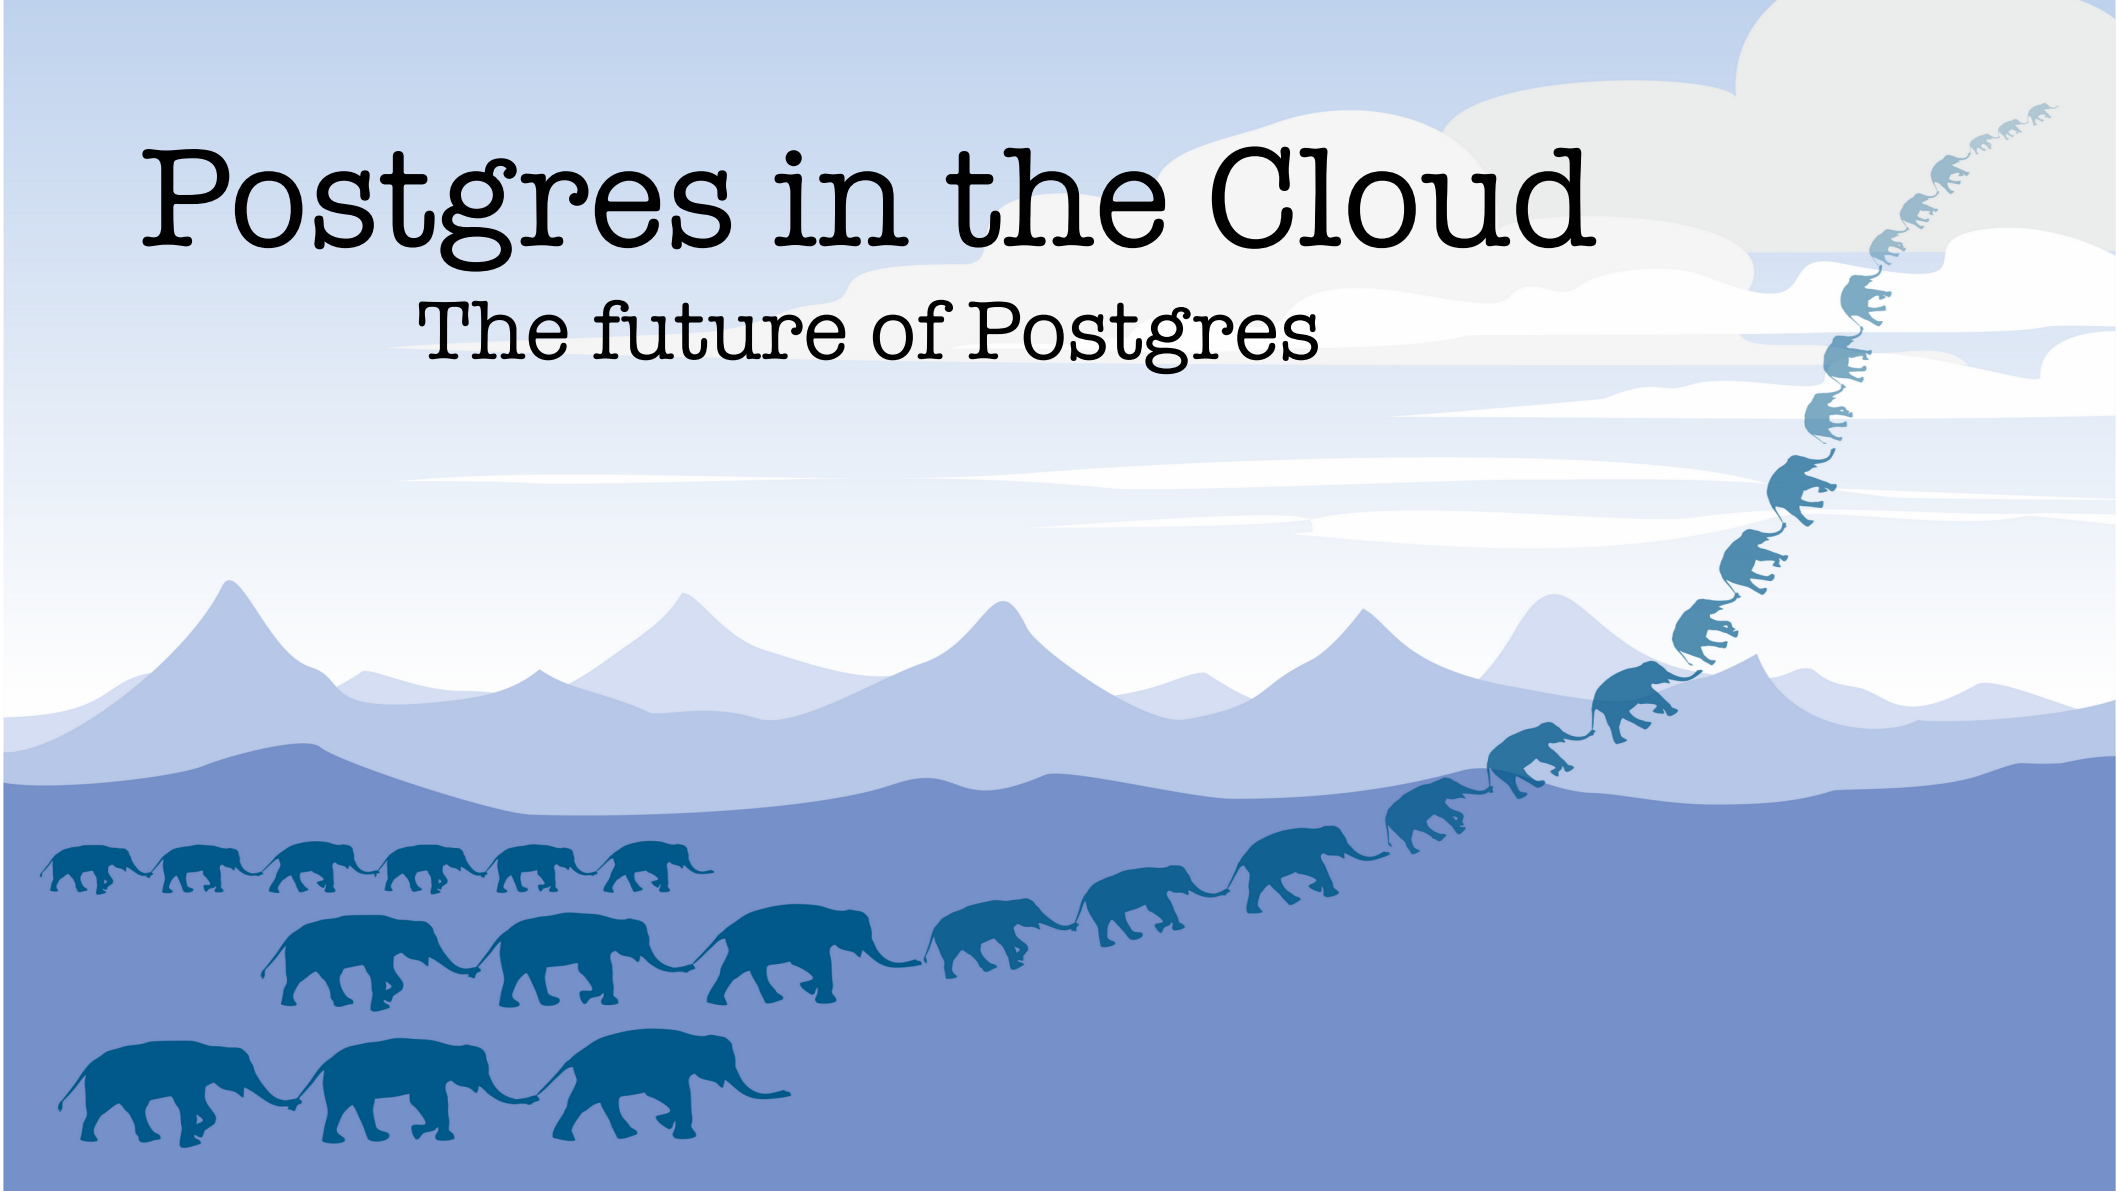

PostgreSQL Future

- Clouds — default platform for databases, 75% of databases in 2022 will be in clouds (Gartner, June 2019)
- Challenges to Postgres
- Zero administration
  Adaptive Postgres, Seamless upgrade (bugfix,security)
  Scalability — Built-in sharding, Numa support, multicore improvements
- Multitenancy
- More storages, better NoSQL, Blockchain support


ALL
YOU
NEED
IS
POSTGRES