



Postgres in a world of DevOps

MOVE FAST, OR ...

MARC LINSTER
SVP, PRODUCT DEVELOPMENT AND SUPPORT

As every company becomes a software company,
DevOps is the new lean manufacturing.



PaaS

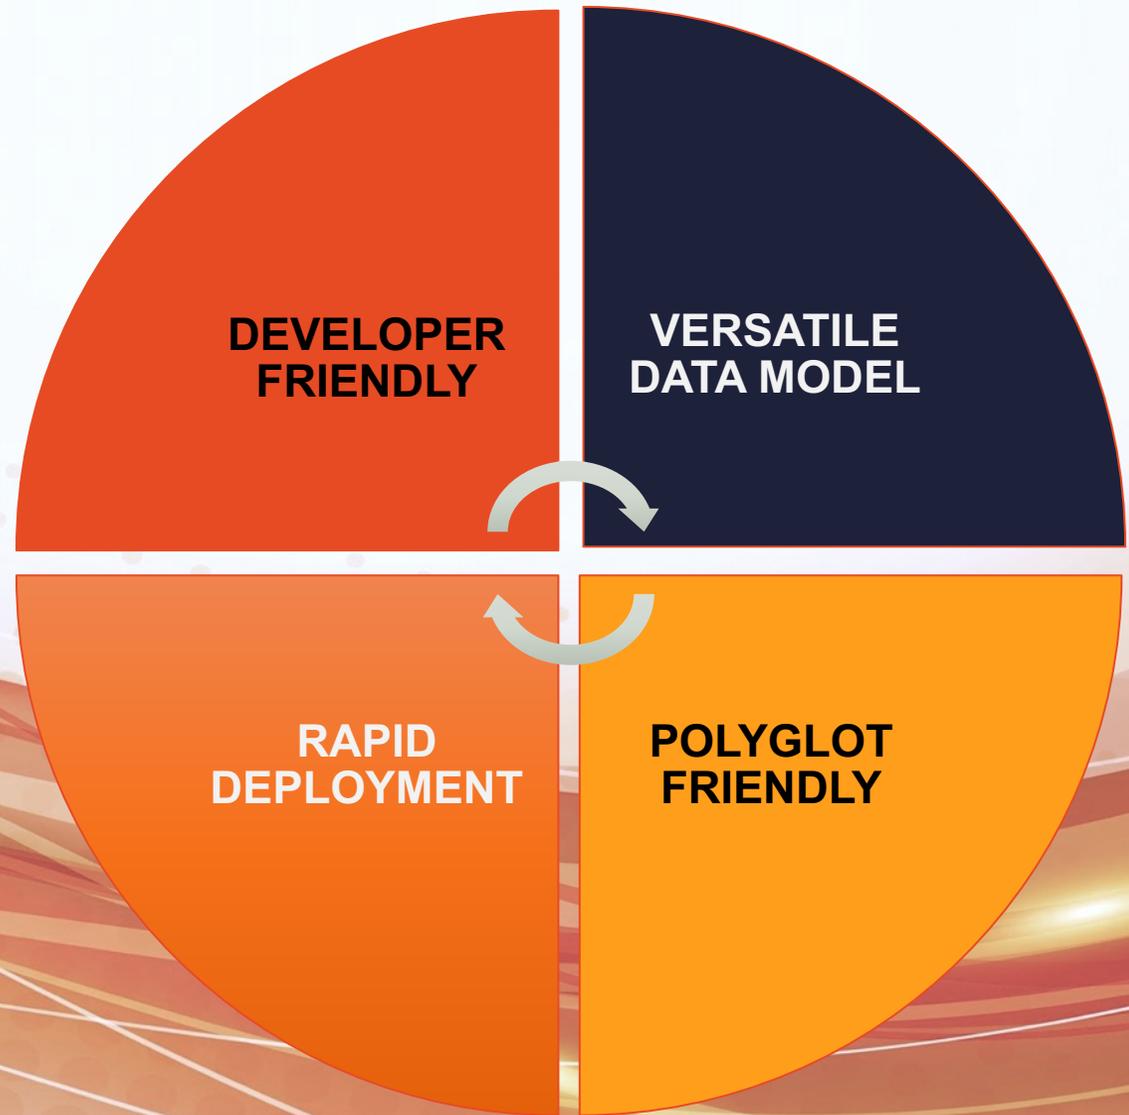
CICD

Microservices

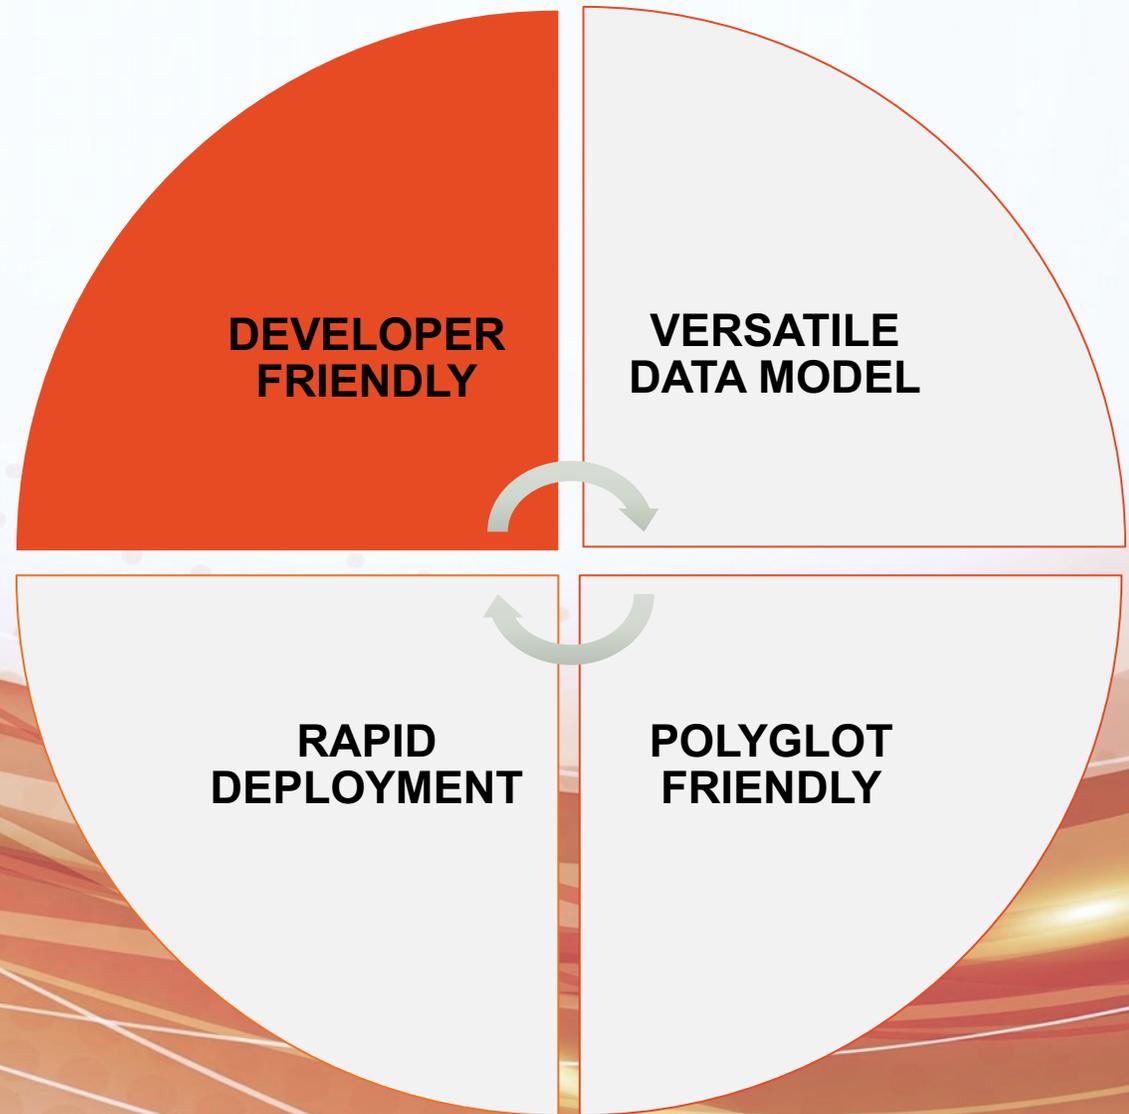
Containers

EDB[™]
POSTGRES

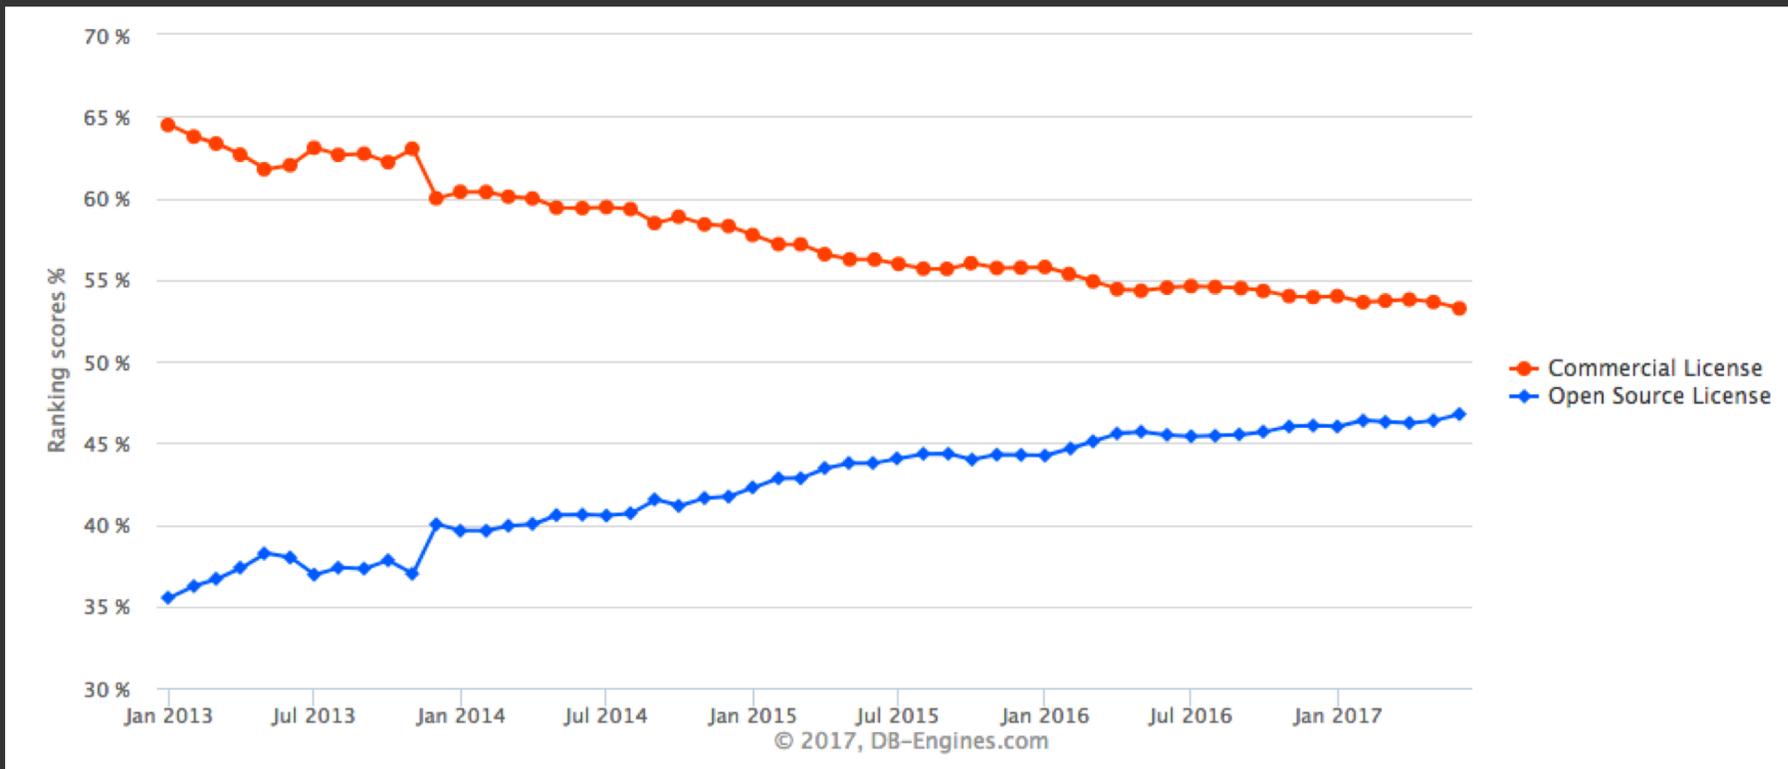
What does it mean for the database?



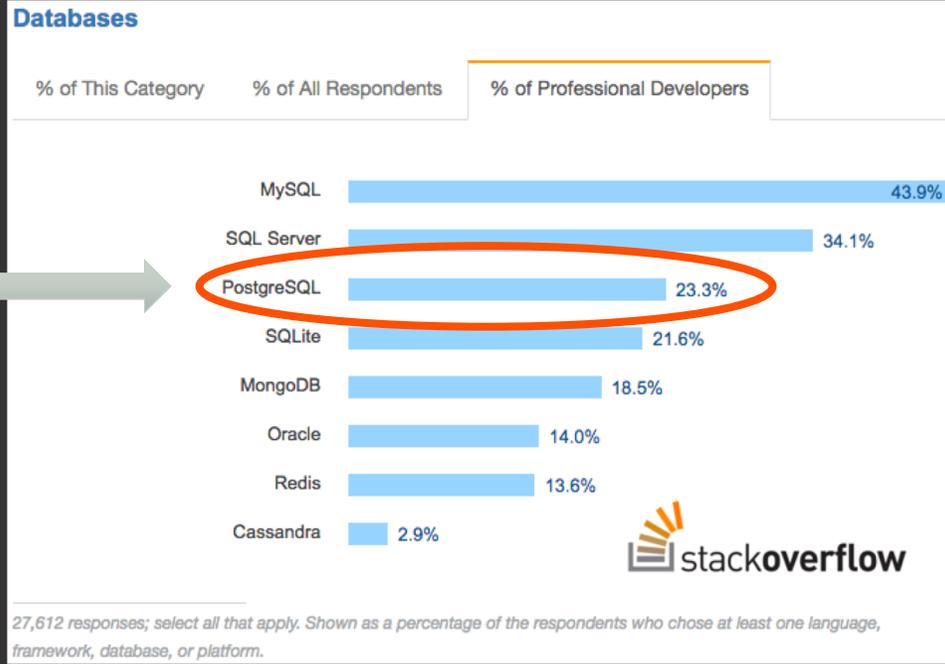
What does it mean for the database?



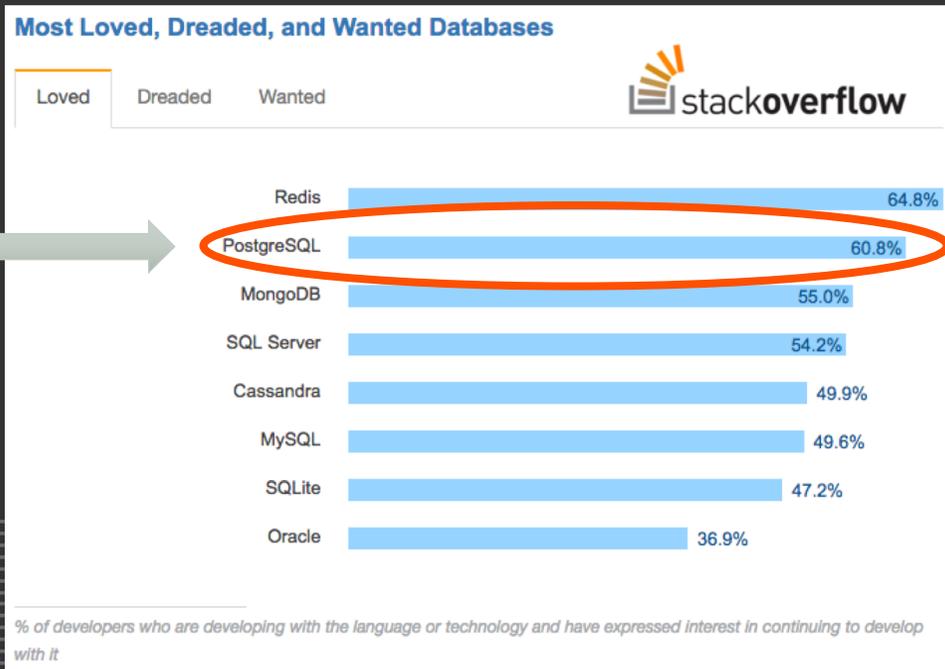
Popularity trend driven by cost and flexibility



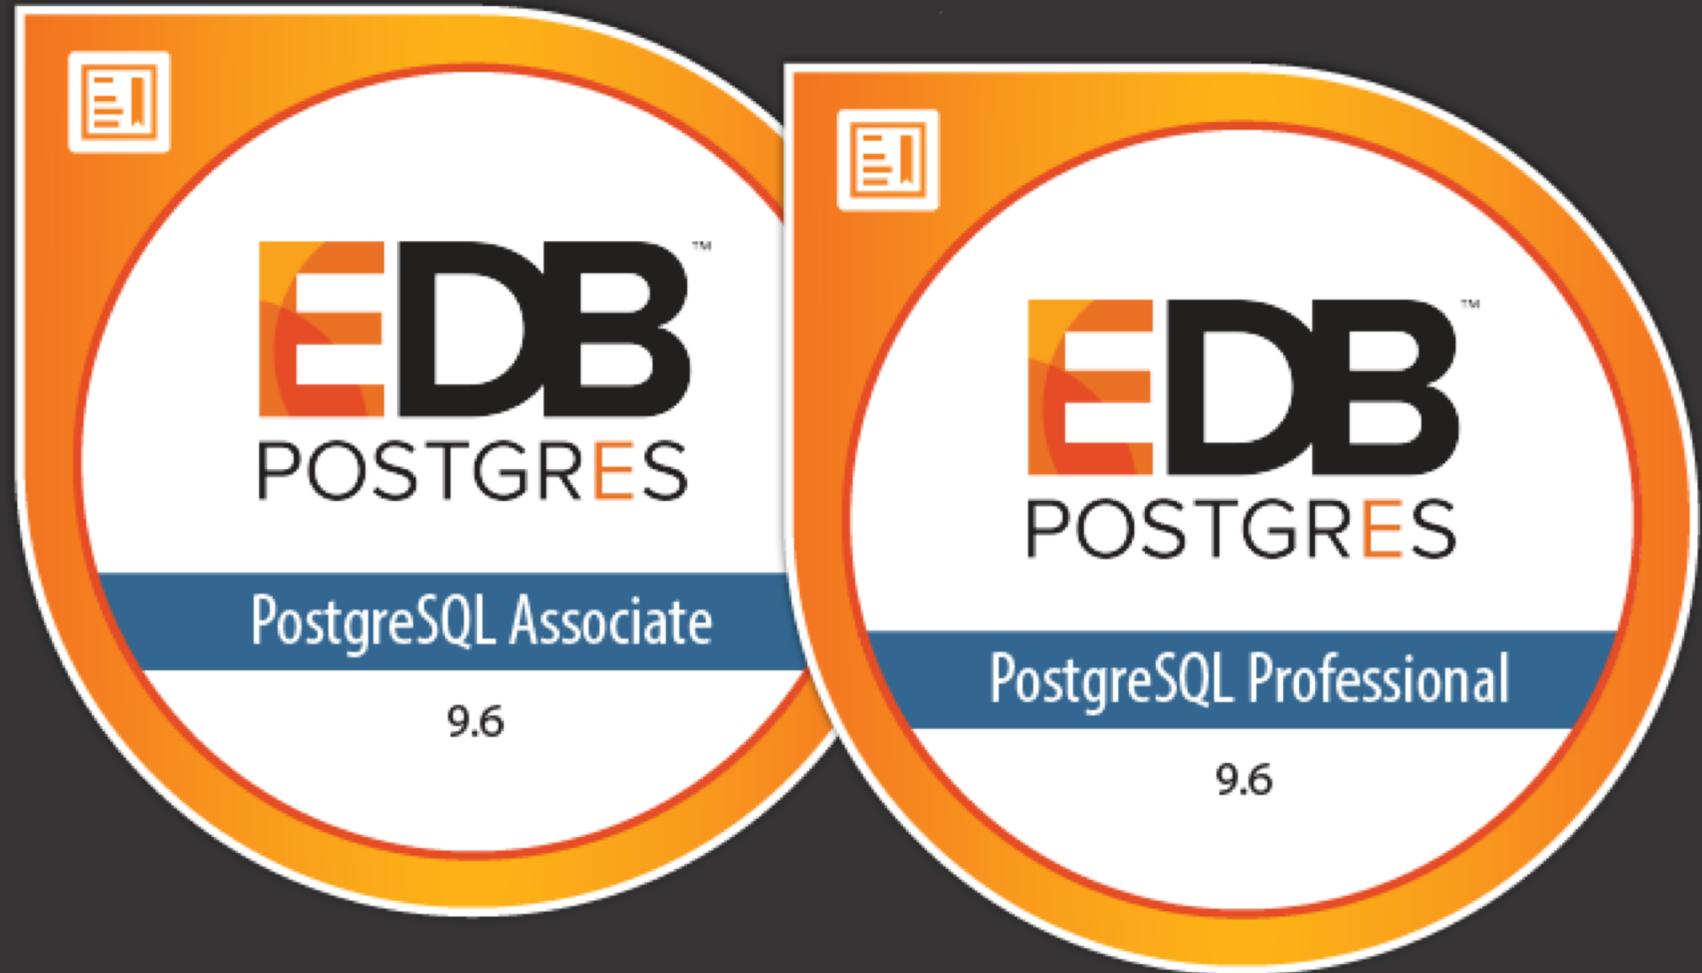
MOST USED



MOST LOVED



Industry recognition of skill proficiency



Postgres Certification



Major DB Developer Tooling

The screenshot displays the Toad Edge Preview interface for a database connection named 'enterprisedb@localhost:5444 (mycompany)'. The main window shows a SQL script with the following content:

```
// (  
78 phone_nbr varchar(10),  
79 visible integer DEFAULT 4,  
80 red_char char default 'x')  
81 RETURNS varchar(10)  
82 /*  
83 Replaces all digits, except for the last visible digits, with the redaction char  
84 */  
85 AS  
86 $$  
87 SELECT overlay (phone_nbr placing rpad(red_char, length(phone_nbr) - visible, red_char) from 1);  
88 $$  
89 LANGUAGE SQL SECURITY DEFINER;  
90  
91 CREATE OR REPLACE VIEW redacteddata.employees AS  
92 SELECT  
93 id,  
94 name,  
95 redact_ssn(ssn) ssn,  
96 redact_phonenbr(phone) phone,  
97 redact_date(birthday) birthday ,  
98 redact_salary(salary) salary,  
99 redact_email(email) email  
100 FROM employeedata.employees;
```

The Object Explorer on the left shows the database structure, including the 'redacteddata' schema and the 'employees' view. The SQL Monitor at the bottom displays the execution results:

Statement	Schema	Date	Count
CREATE OR REPLACE VIEW redacteddata.employees AS SELECT id, name, redact_ssn(ssn) ssn, redact_phonenbr(phone) phone, redact_date(birthday) birthday, redact_salary(salary) salary, redact_email(email) email FROM employeedata.employees;	public	Apr 15, 2018 11:37:59 AM	2
CREATE OR REPLACE FUNCTION redact_date (input_date date) RETURNS date /* sets the year to 0 */ AS \$\$ SE public	public	Apr 15, 2018 11:37:47 AM	1
CREATE OR REPLACE FUNCTION redact_phonenbr (phone_nbr varchar(10), visible integer DEFAULT 4, red_cha public	public	Apr 15, 2018 11:35:19 AM	1
CREATE OR REPLACE FUNCTION redact_salary (salary money) RETURNS money /* always returns 0 */ AS \$\$ SE public	public	Apr 15, 2018 11:34:57 AM	1
CREATE or REPLACE FUNCTION redact_email (email varchar(100), visible integer DEFAULT 1, red_char char def public	public	Apr 15, 2018 11:34:13 AM	1
CREATE OR REPLACE FUNCTION redact_ssn (ssn varchar(11)) RETURNS varchar(11) /* replaces 020-12-9876 wi public	public	Apr 15, 2018 11:33:15 AM	2

The status bar at the bottom indicates 'AutoCommit ON', 'Writable', 'Insert', and '91 : 9'.

DATABASE WORLD IS CHANGING

DBA and Developer are no longer living in separate worlds

PAST

“the corporate standard”

Company data centers

Waterfall

IT provisioning

Develop –then- operate

Manual tasks

Routine tasks

Database expert

PRESENT

Many new databases

The cloud

Agile

Self-service provisioning

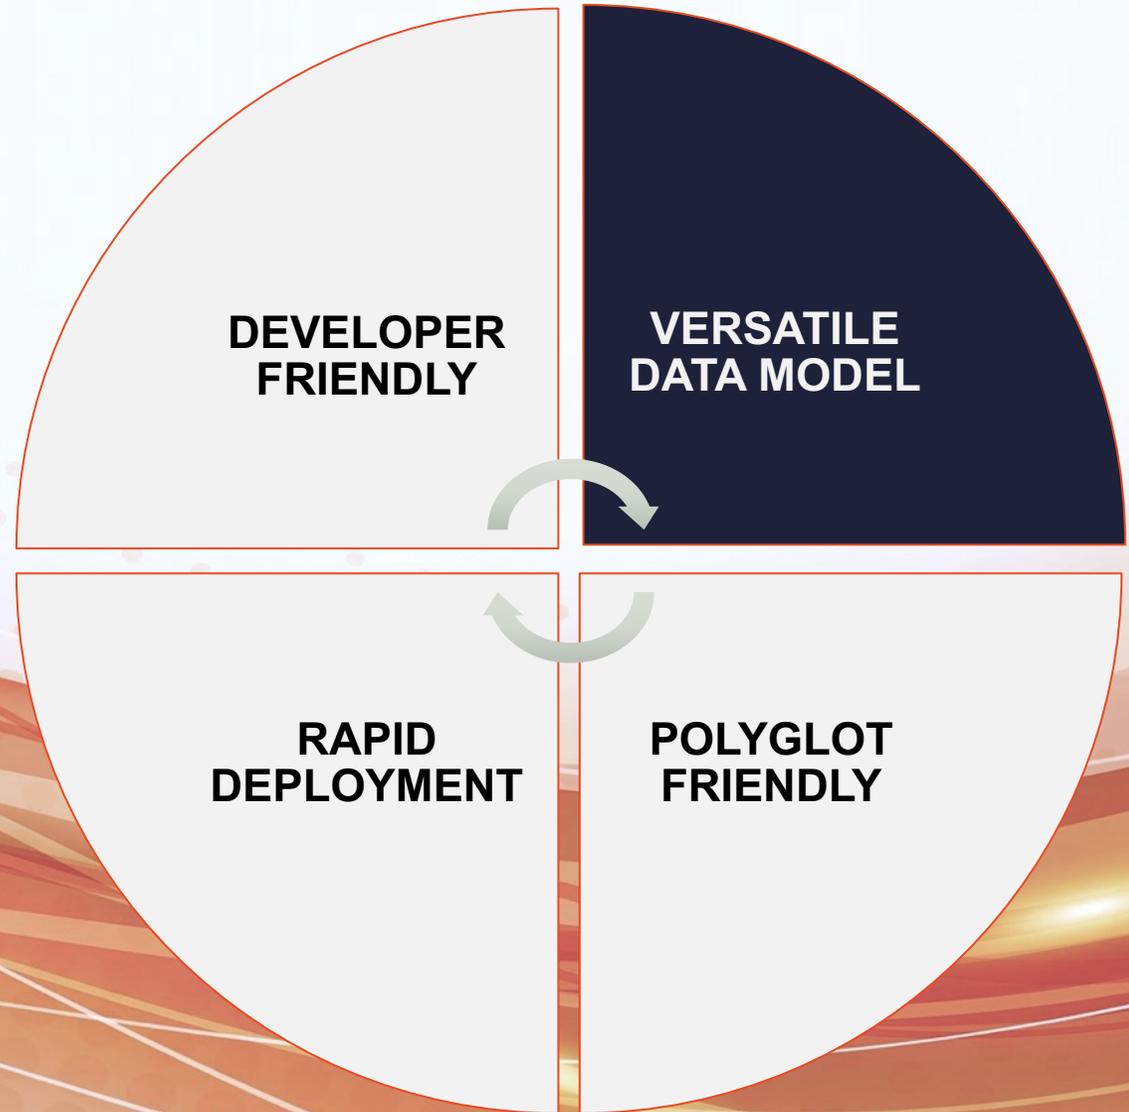
DevOps

Automation

Hard problems

Data Expert

What does it mean for the database?



Postgres: The most versatile DBMS

Runs on all platforms

Speaks every language

Not only SQL

Scales in both directions



JSON Data Types

1. Number:

- Signed decimal number that may contain a fractional part and may use exponential notation.
- No distinction between integer and floating-point

2. String

- A sequence of zero or more Unicode characters.
- Strings are delimited with double-quotation mark
- Supports a backslash escaping syntax.

3. Boolean

- Either of the values true or false.

4. Array

- An ordered list of zero or more values,
- Each values may be of any type.
- Arrays use square bracket notation with elements being comma-separated.

5. Object

- An unordered associative array (name/value pairs).
- Objects are delimited with curly brackets
- Commas to separate each pair
- Each pair the colon ':' character separates the key or name from its value.
- All keys must be strings and should be distinct from each other within that object.

6. null

- An empty value, using the word null

JSON is defined per RFC – 7159
For more detail please refer
<http://tools.ietf.org/html/rfc7159>

JSON Data Type Example

```
{
  "firstName": "John",           -- String Type
  "lastName": "Smith",          -- String Type
  "isAlive": true,               -- Boolean Type
  "age": 25,                     -- Number Type
  "height_cm": 167.6,           -- Number Type
  "address": {                  -- Object Type
    "streetAddress": "21 2nd Street",
    "city": "New York",
    "state": "NY",
    "postalCode": "10021-3100"
  }
  "phoneNumbers": [            -- Object Array
    {                            -- Object
      "type": "home",
      "number": "212 555-1234"
    },
    {
      "type": "office",
      "number": "646 555-4567"
    }
  ],
  "children": [],
  "spouse": null                -- Null
}
```

Why choose between NoSQL and Relational?

ANSI SQL

```
SELECT DISTINCT
  product_type,
  data->>'brand' as Brand,
  data->>'available' as Availability
FROM json_data
JOIN products
ON (products.product_type=json_data.data->>'name')
WHERE json_data.data->>'available'=true;
```

JSON

product_type	brand	availability
AC3 Phone	ACME	true

START SCHEMALESS



- Leverage structure as it emerges
- Support agile, iterative development
- Create data models where they provide value

LEVERAGE JSONB

- Leverage JSONB for rapidly changing data models
- Example: address data records
 - Conventional columns: First Name, Last Name
 - JSONB: Contact Information
 - Phone numbers (home, cell, car, weekend, boyfriend...)
 - Email (work, private, spam ...)

AVOID PENALTY



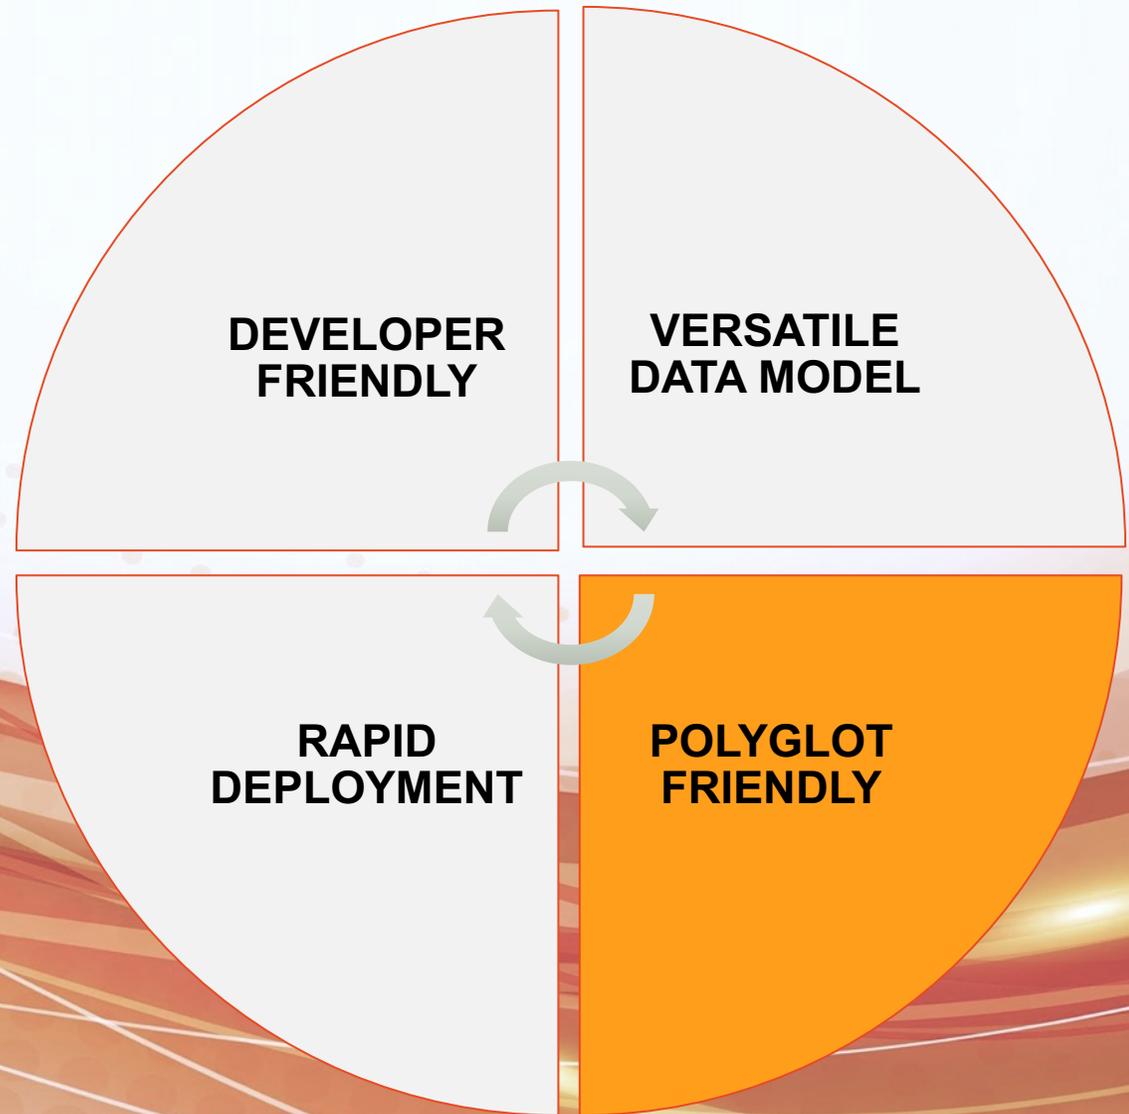
- Avoid the DDL penalty for adding columns



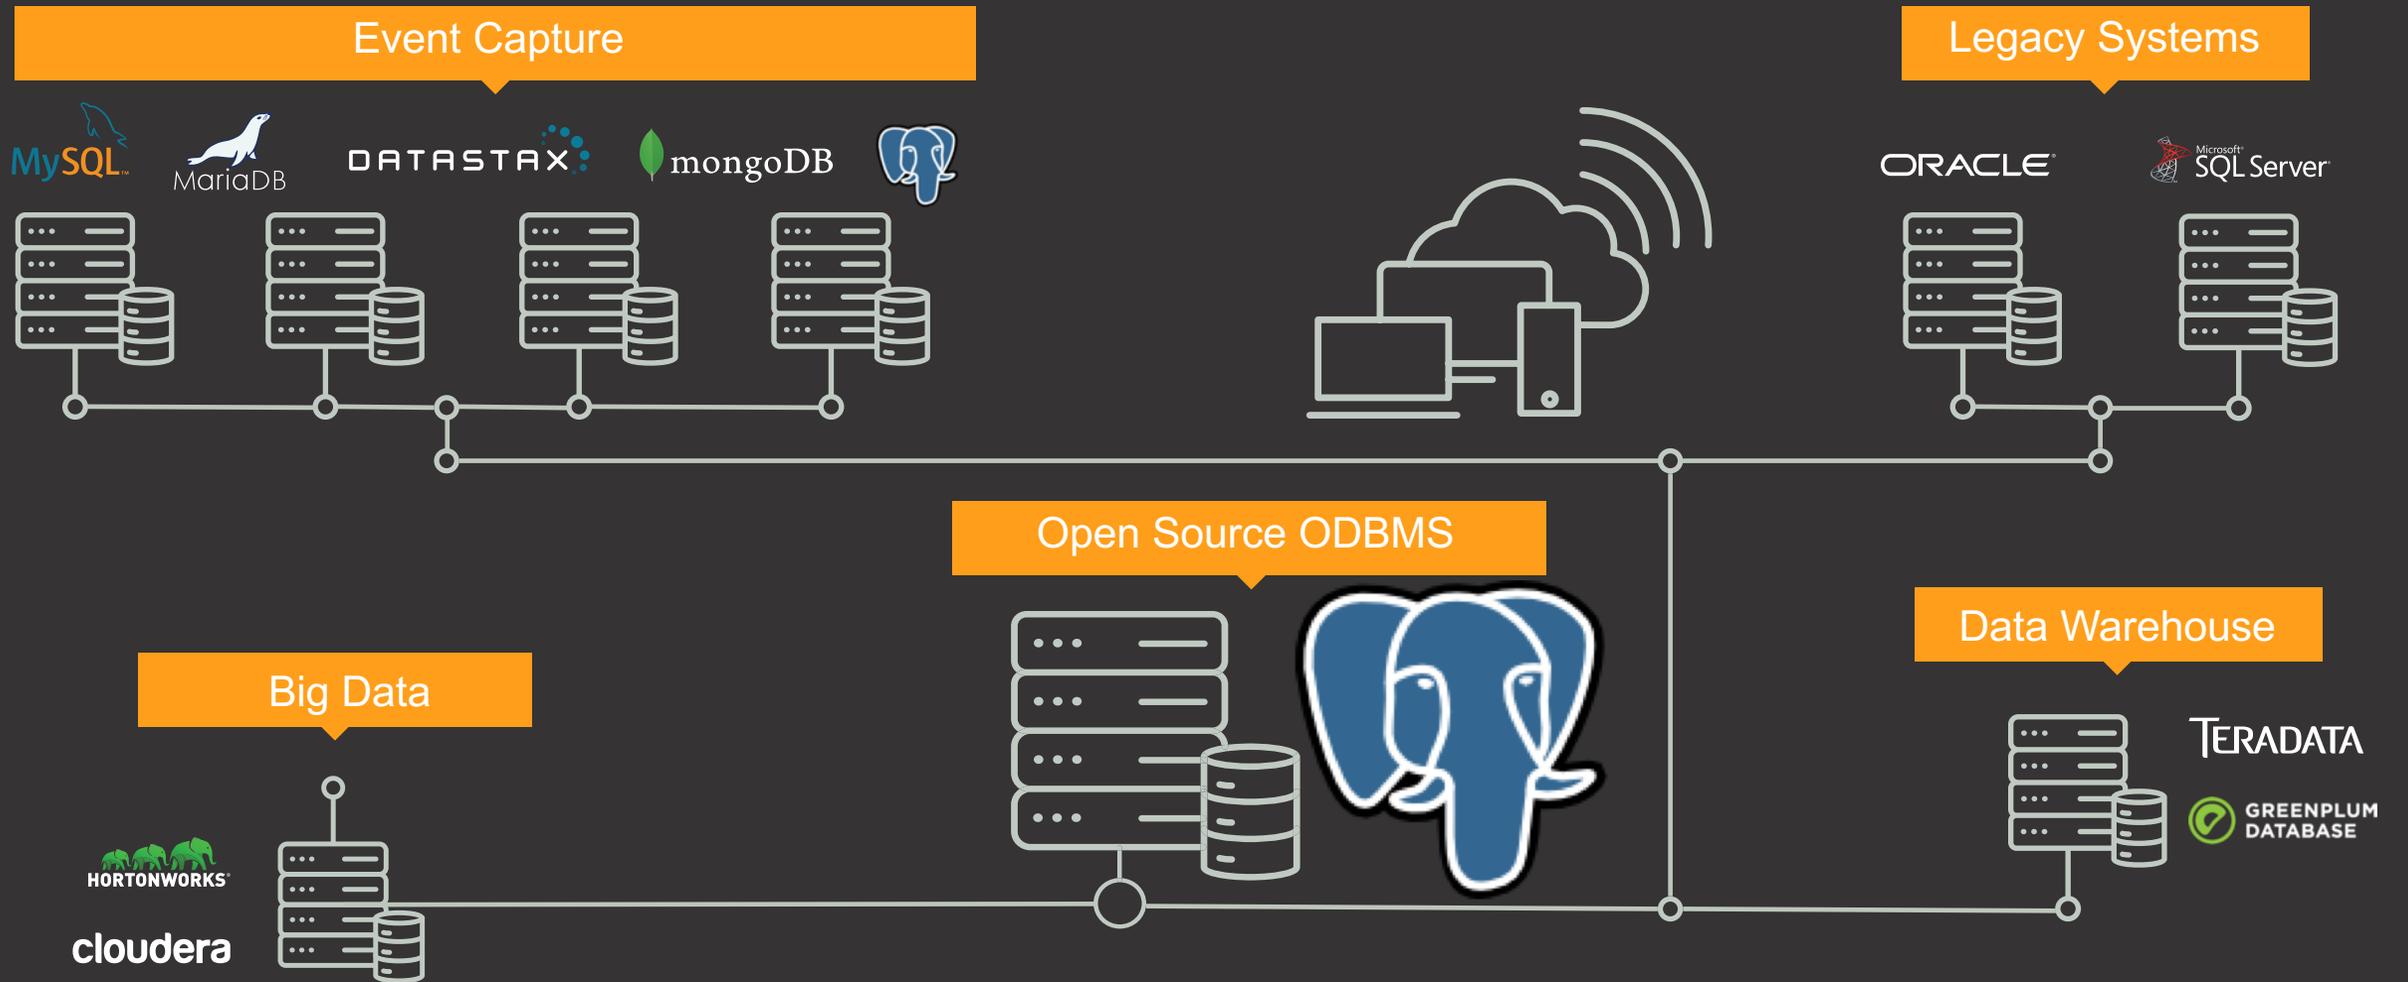
Why not just SQL?

EDB
POSTGRES

What does it mean for the database?



Postgres in the Polyglot Ecosystem



FDW IMPLEMENTS SQL/MED ("SQL MANAGEMENT OF EXTERNAL DATA")

PostgreSQL 9.1 - read-only support

PostgreSQL 9.3 – read/write support

PostgreSQL 9.6 – pushdown joins, sorts, UPDATE, DELETE

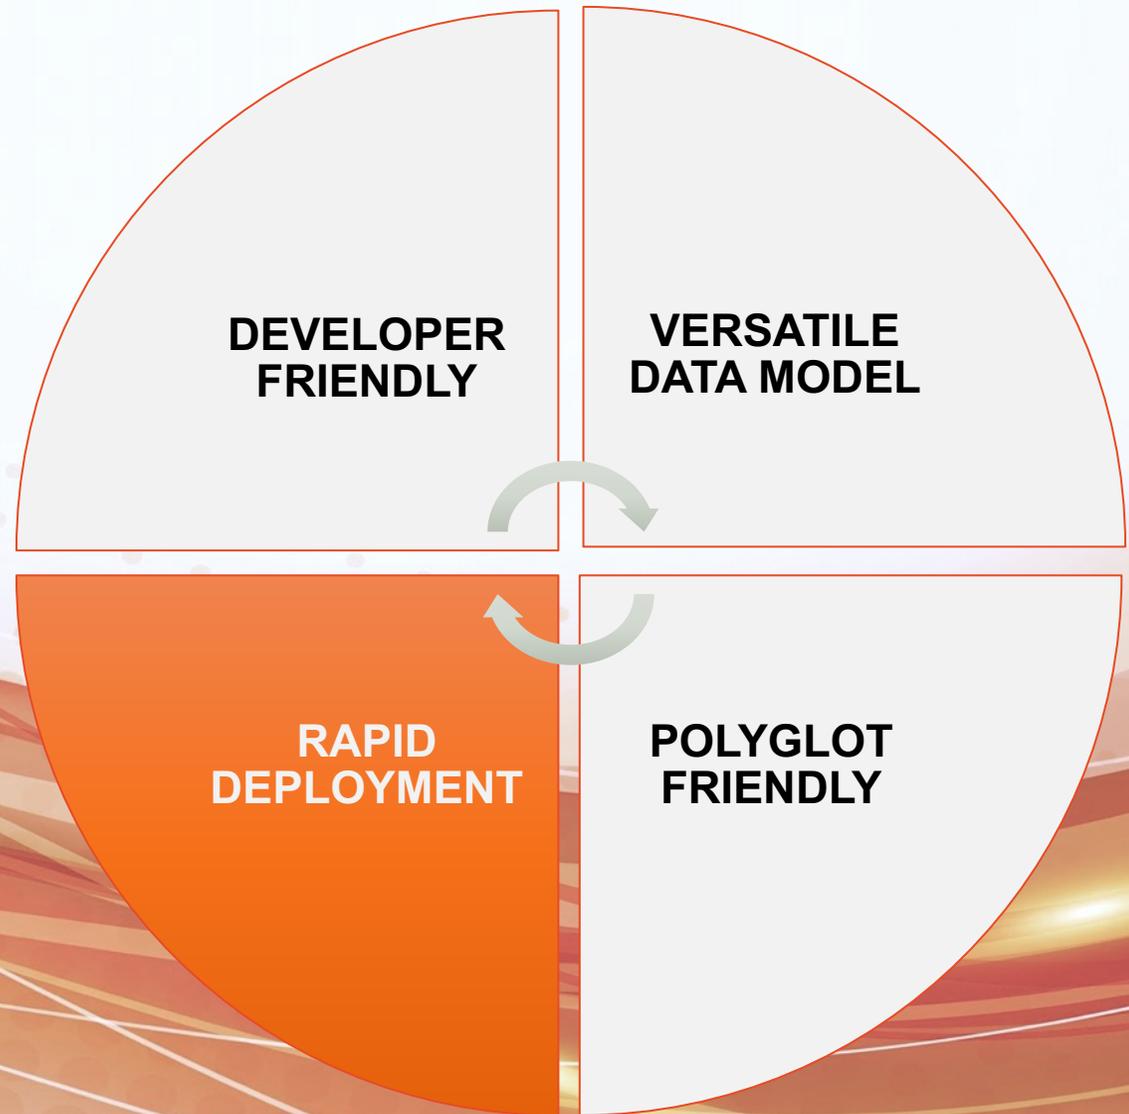
PostgreSQL 10 - aggregate pushdown

*FDW: Makes data on other servers (or services) look like tables in Postgres.
Available for many data sources (MongoDB, MySQL, HDFS, Spark, ...)*



Foreign Data Wrappers – Making Polyglot Happen

What does it mean for the database?





CLOUD

- Public (AWS, Google, Azure, Alibaba Cloud, ...)
- Private (OpenStack, VCloud, Puppet/Chef on Virtualization, Pivotal Cloud Foundry, Container/Kubernetes/OpenShift)



DBaaS

- Stop deploying databases
 - Deploy clusters w. HA, DR, self healing, scaling, etc.
 - provide services, not databases



Micro Services

- Large monolithic databases refactored into transaction sets
- Applications get refactored into micro-services
- Deployment models move from VM/Bare Metal to DBaaS and Containers



Rapid Deployment

EDB[™]
POSTGRES

SUPPORTING DevOps WITH DBaaS



Address the tension between developers and operations

DEVELOPERS WANT | Agility | Speed to deploy |
Flow through tool chain

OPERATIONS WANTS | Visibility | Control –
budgeted resources – cost – data models |
Efficient use of resources



Micro Services and Containers

MAJOR TRENDS

- Large monolithic databases refactored into transaction sets
- Applications get refactored into micro-services
- Deployment models move from VM/Bare Metal to DBaaS and Containers



**Monolithic DB
System of Record**

*Refactored by
transaction
sets*



Shipping
Transactions



Payroll
Transactions



Inventory
Transactions



Salary
Transactions



Sales
Transactions



Compensation
Transactions



Customer
Transactions



Employee
Transactions



Key to DevOps

EDB
POSTGRES

THE CHANGING ROLE OF THE DBA

From Ops Supporter to DevOps Collaborator



1

FROM DATABASE ADMIN TO DATA STRATEGIST

Automation ends mundane and repetitive tasks

Adding new servers by clicking a few buttons

Frees up time for strategic efforts

More focus on the data and applications

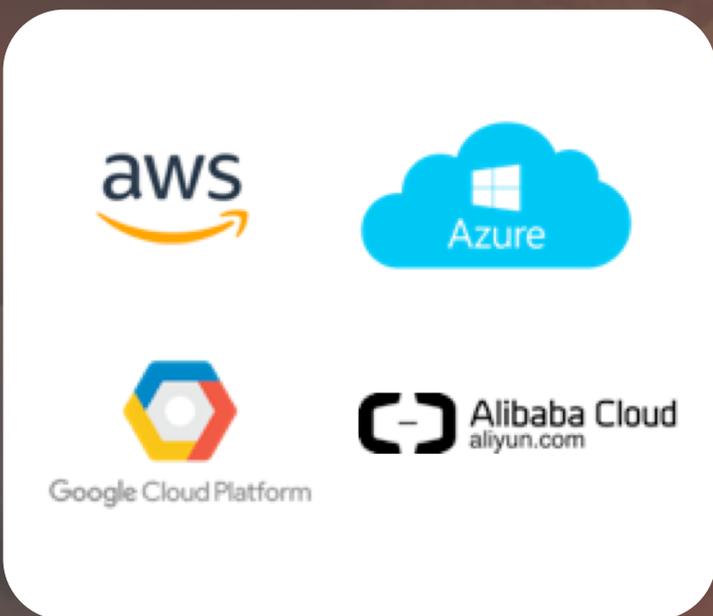
Understanding data sources and value to organization

Develop new skills in areas such as data science and machine learning

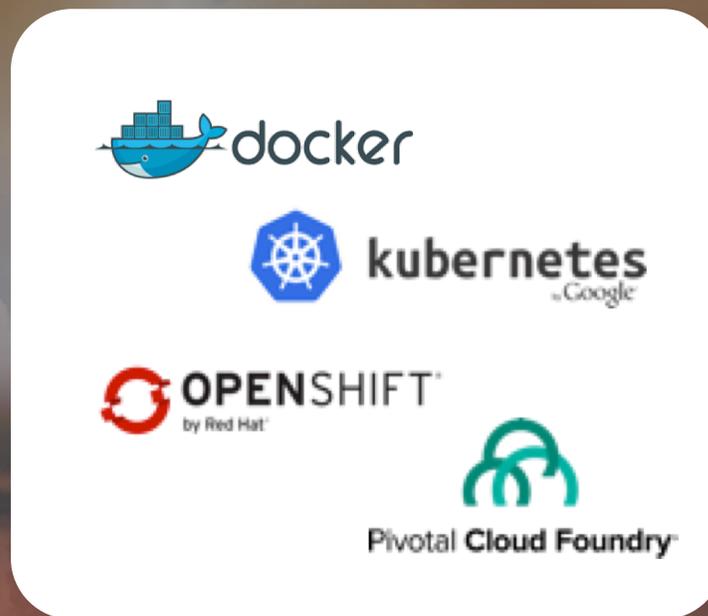
2

FROM HARDWARE TWEAKER TO MASTER OF CLOUD TECHNOLOGIES

CLOUDS

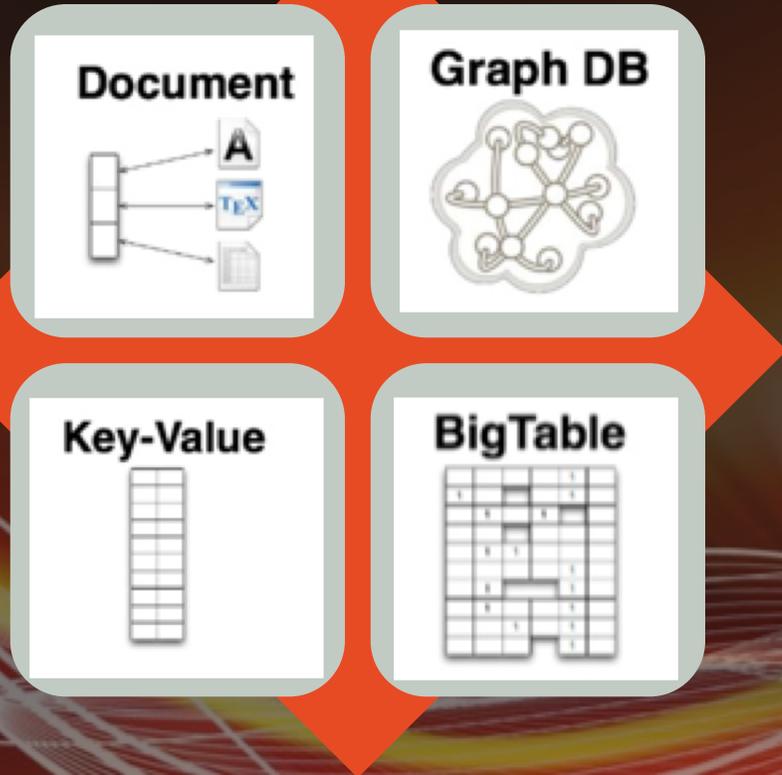


CONTAINERS



3

FROM SCHEMA INDEXER TO MASTER MODELER



DBAs historically specialized by RDBMS

Now need to know/recommend different platforms

Also still need deep DB knowledge

And understand DB as part of a full stack

4

FROM SERVER JOCKEY TO UTILITY PROVIDER

DBs now viewed as a utility

Flip a switch and get it now

Burstable capacity

Less time to plan

5

FROM LICENSE GATEKEEPER TO BUDGET OWNER

Most cloud services are provided by a third-party vendor

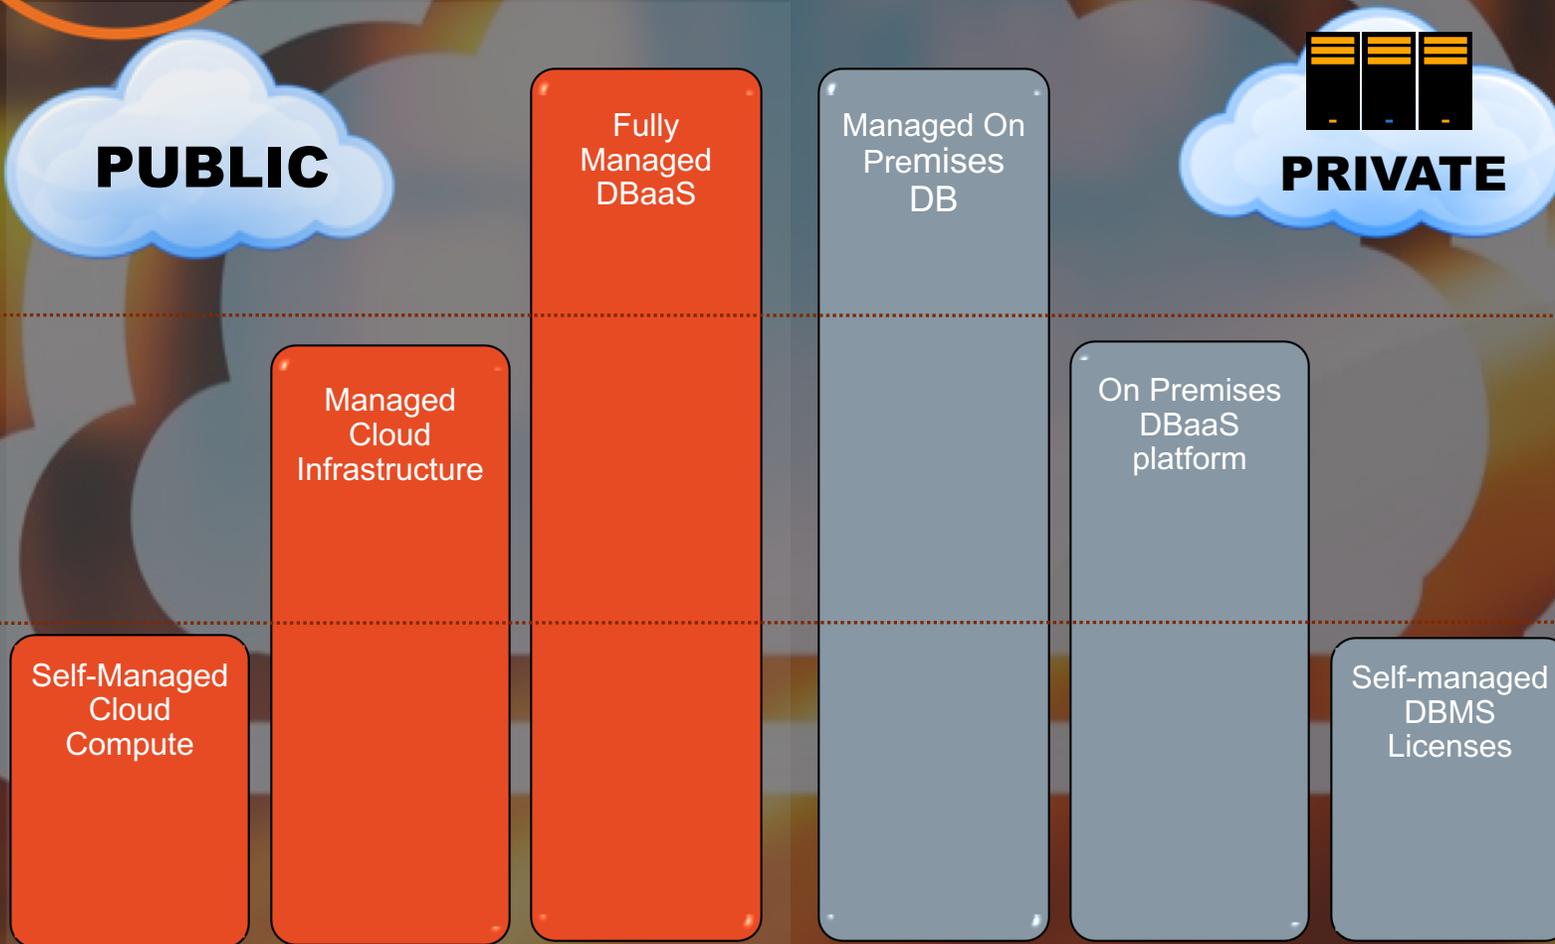
Need to understand services

How to optimize cost vs. performance including infrastructure

Architecting to meet requirements and budget

6

FROM SITE ADMIN TO PLATFORM BROKER



DBs becoming platform agnostic

Support on premises or public cloud

Many deployment and service level options

Collaborate with teams supporting the cloud/full application stack

7

FROM QUERY OPTIMIZER TO PERFORMANCE PRO

Less time updating/troubleshooting servers

More time optimizing data architecture

Monitoring of infrastructure and database performance holistically

Advanced data modeling, query tuning and Indexing strategies to improve performance

8

FROM TO ACCOUNT ADMIN TO SECURITY EXPERT

Cloud vendors provide secure infrastructure

But must ensure that all systems are using that infrastructure properly, especially the database

Protecting sensitive information across highly-distributed infrastructures can be a challenge

DBA needs to understand possible threat sources

Compliance with corporate, industry, and governmental regulations add new responsibilities on DBA



9

FROM OPS SUPPORT TO DevOps COLLABORATOR

Increasing focus on continuous deployment and delivery

DBAs take on tasks once reserved for developers and work closely with DevOps teams

More moving out of centralized IT and into the line-of-business/applications teams

While continuing as data steward and trusted adviser

10

AND STILL NEED TO KEEP IT RUNNING

Service providers apply fix packs and upgrade DBMS versions

Backup and recovery in the cloud is fully automated

DBA still needs awareness and understanding of the potential impact

NOW HAVE TO WORRY ABOUT THE CLOUD FOR:

Scalability

Fault tolerance

Replication

**DBA IS STILL
ULTIMATELY RESPONSIBLE**

DATABASE AS A SERVICE

DBaaS CAN HELP

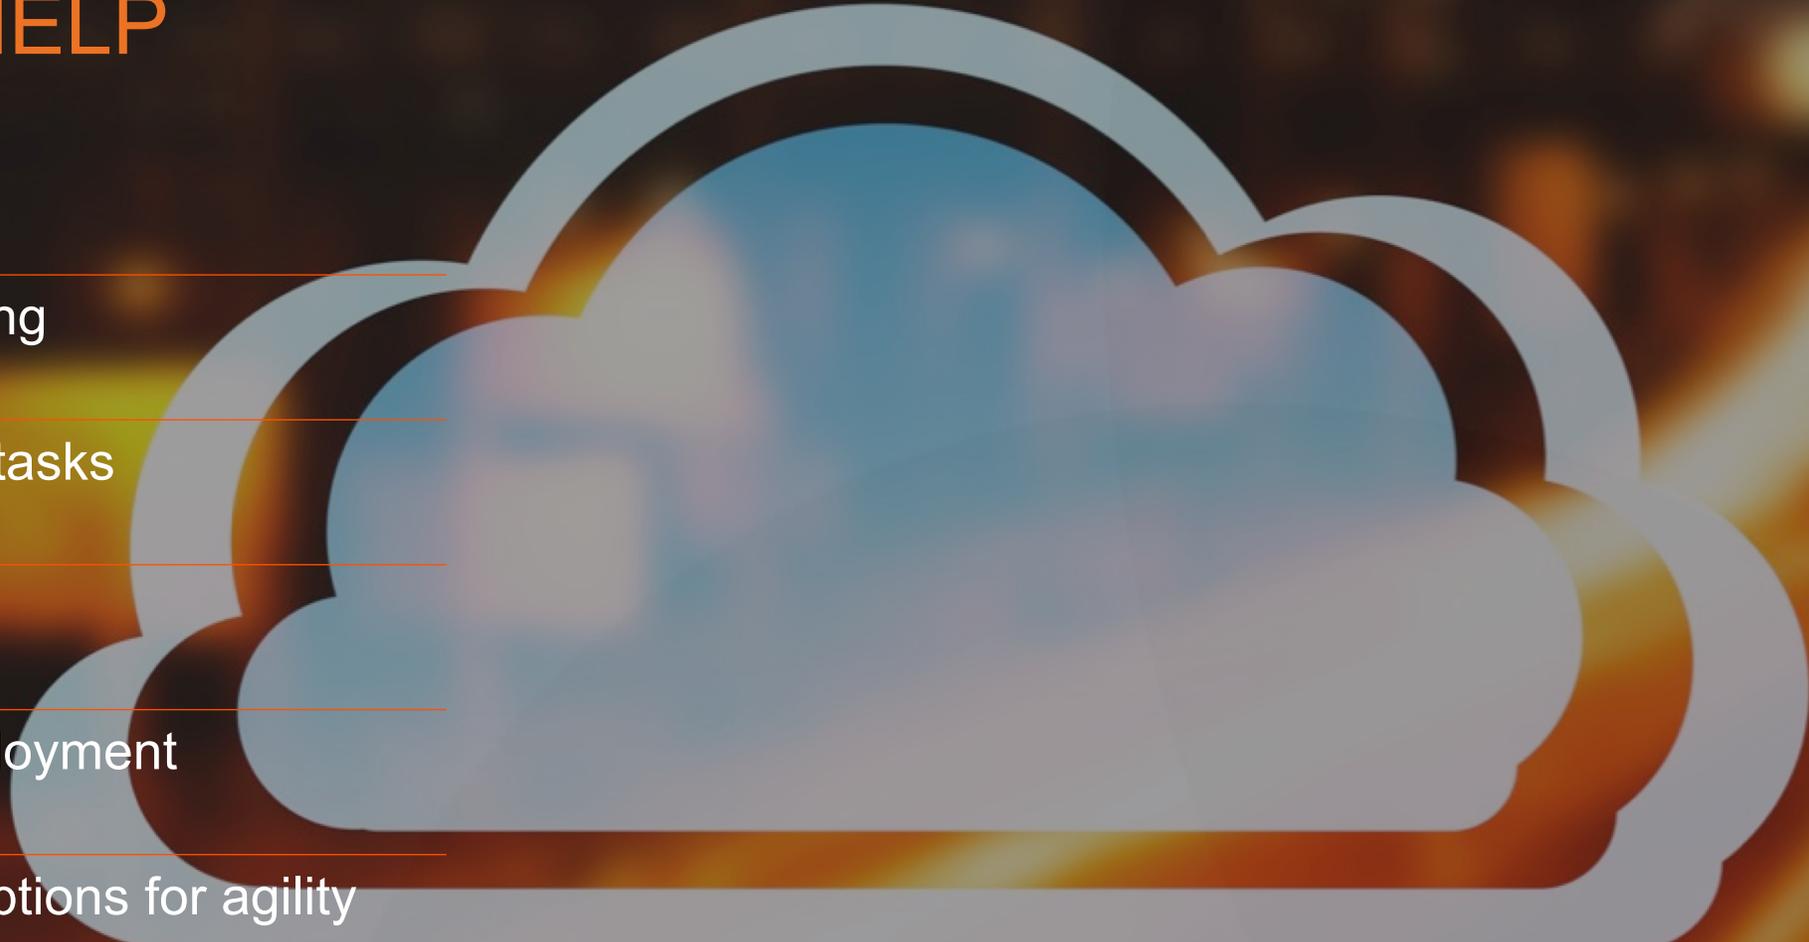
Self-service provisioning

Automation of routine tasks

“Infinite” capacity

Multi-cloud/hybrid deployment

Tradeoff control and options for agility



What does it mean for the database?

- Easy to use
- High adoption rate
- Readily available skill set
- Training and certification

DEVELOPER FRIENDLY

- Multi model database
- Rich set of data types and extensions (JSONB, Hstore, PostGIS, ...)

VERSATILE DATA MODEL

RAPID DEPLOYMENT

- Cloud friendly
- Micro/Mini services database refactoring
- DBaaS focus

POLYGLOT FRIENDLY

- Open co-existence with HDFS, Mongo, Kafka, ...
- Rich FDW library



THANK YOU