EDB Postgres Database Containers

DATABASE PLATFORM

FOR DIGITAL BUSINESS

Ossi Karjalainen

Director, Sales Engineering EMEA

Digital Business - A Rapidly Changing Universe

Traditional	DBaaS	DevOps

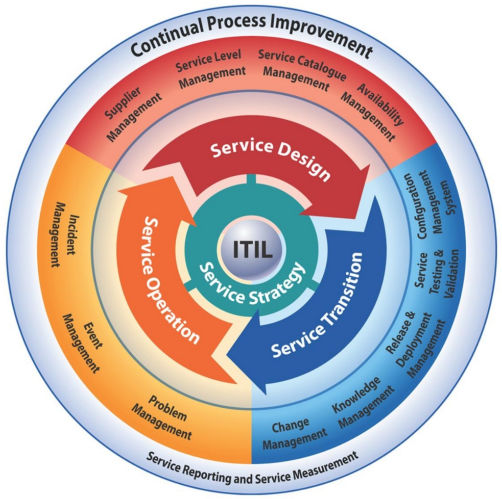

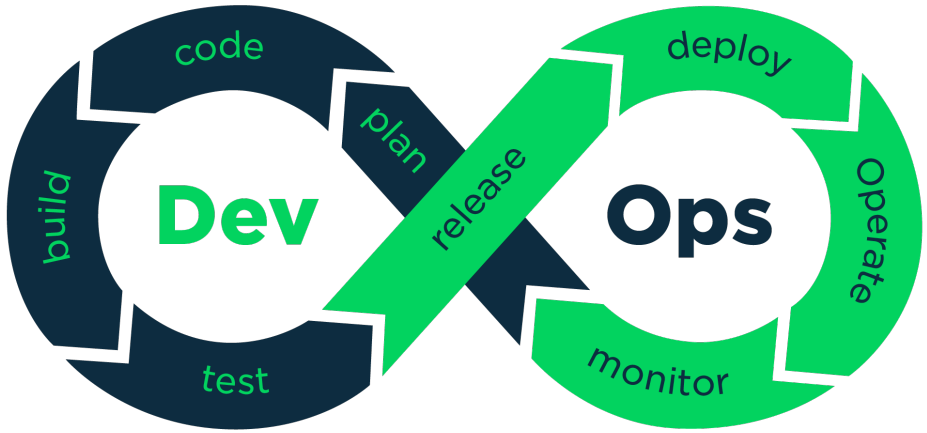

Polyglot Multi Model Data World

Database Platform for Digital Business

- Everybody is moving towards Digital Business
  - Customer engagement
  - Self services
  - New value-add
  - Lower cost
  - Greater flexibility

EDB POSTGRES

Digital Business - Enterprise Requirements

- Databases must be extremely reliable, highly available and resilient
- Solutions need to leverage the legacy through integration or migration
- Rapid innovation is the name of the game

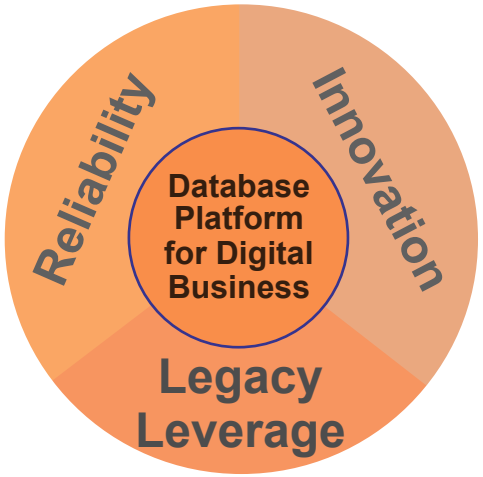

Digital Business - Enterprise Requirements

- Databases must be extremely reliable, highly available and resilient
- Solutions need to leverage the legacy through integration or migration
- Rapid innovation is the name of the game

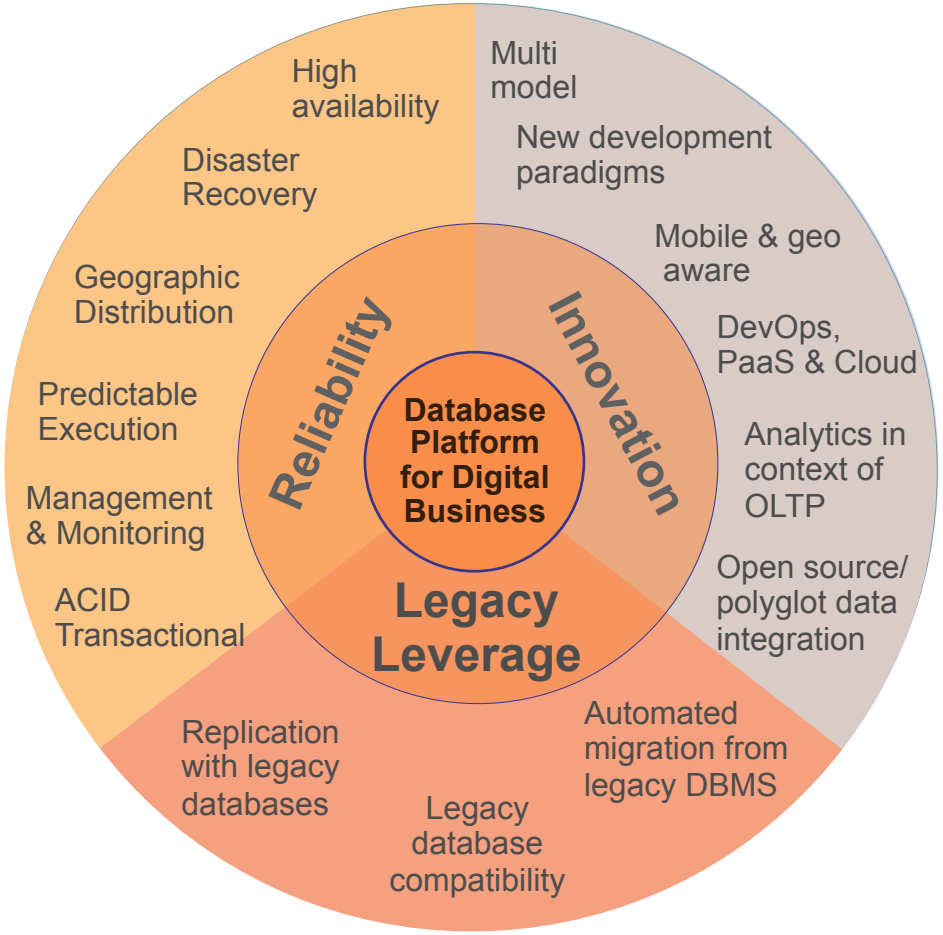

Enterprise Requirements – Focus on Innovation

- Databases must be extremely reliable, highly available and resilient
- Solutions need to leverage the legacy through integration or migration
- Rapid innovation is the name of the game

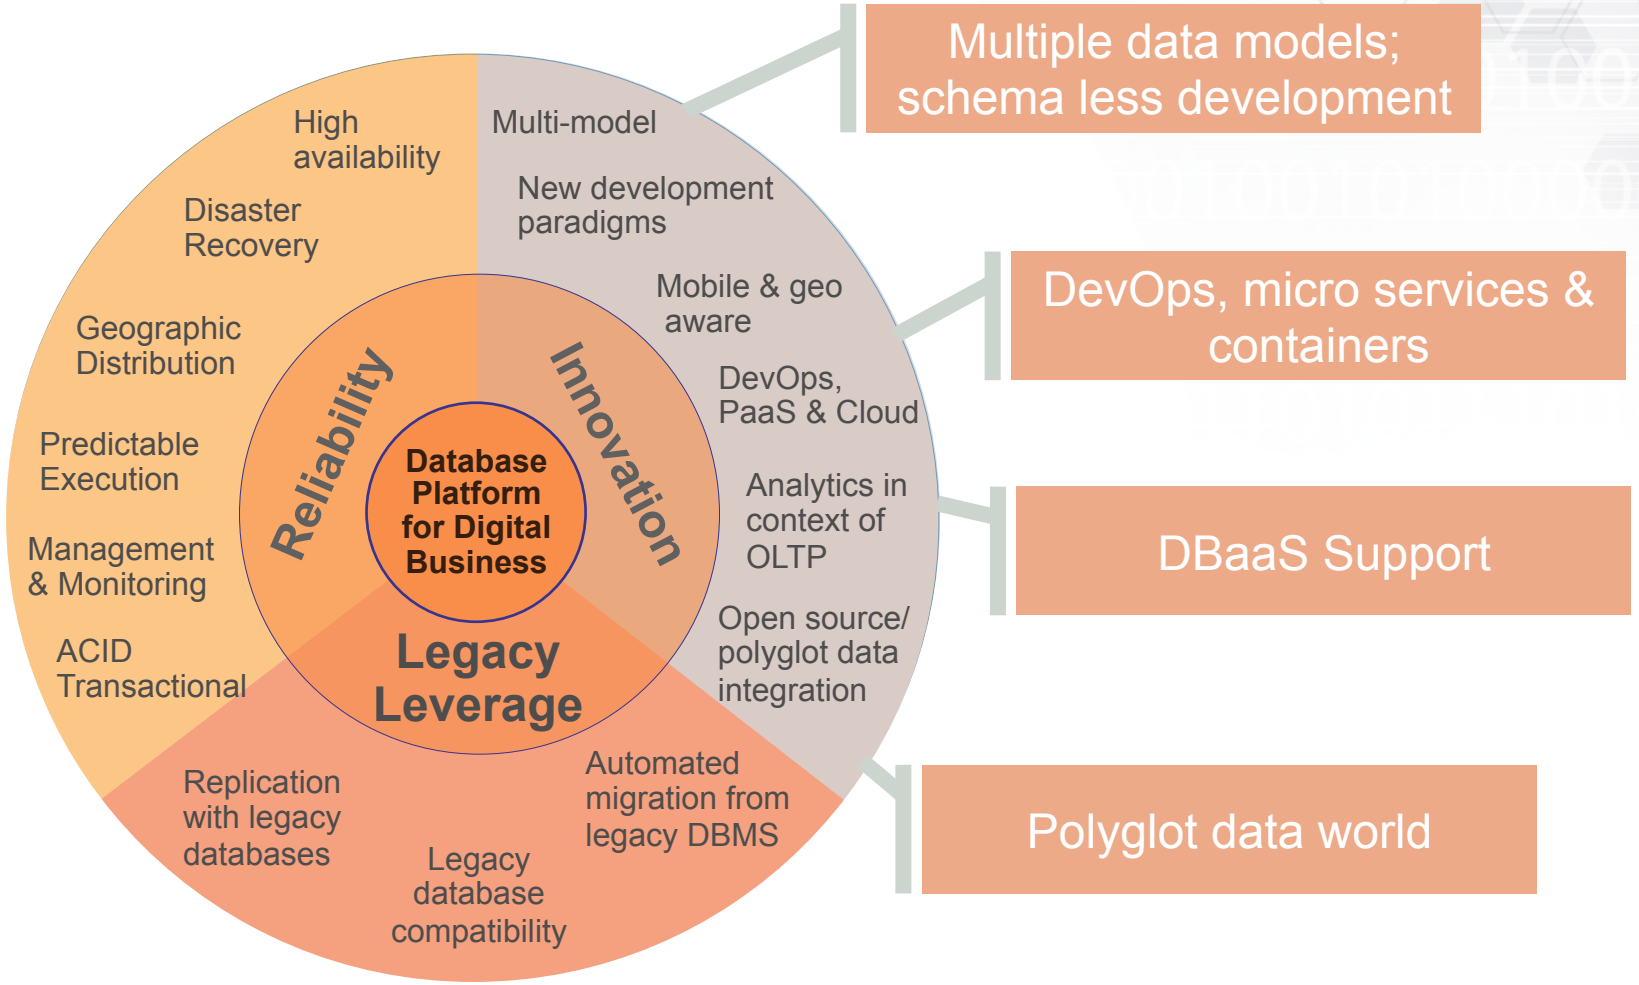

Focus on Innovation – Multi Model for the Enterprise

- Data Types
  - Postgres has JSON, JSONB, Key-Value Pair, plus arrays, ranges, timezones, dates, integer, floating point, etc.
- Performance Benchmarks
  - Postgres is very fast and can handle huge amounts of data
  - Postgres can selectively relax key ACID features to increase performance
- Flexibility
  - Start schema less – support DevOps – 'We need no DBA'
  - Leverage structure as it emerges
  - Manage transactional, document, IoT, … in one environment

EDB POSTGRES

Focus on Innovation – Polyglot Data World

- It's a fact not one DBMS technology dominates
- Event capture or caching of transactions on the edge
- Analysis of large data sets in the Hadoop ecosystem
- Transactional systems are still at the heart
- Existing systems have to be leveraged rather than replaced in order to move forward efficiently

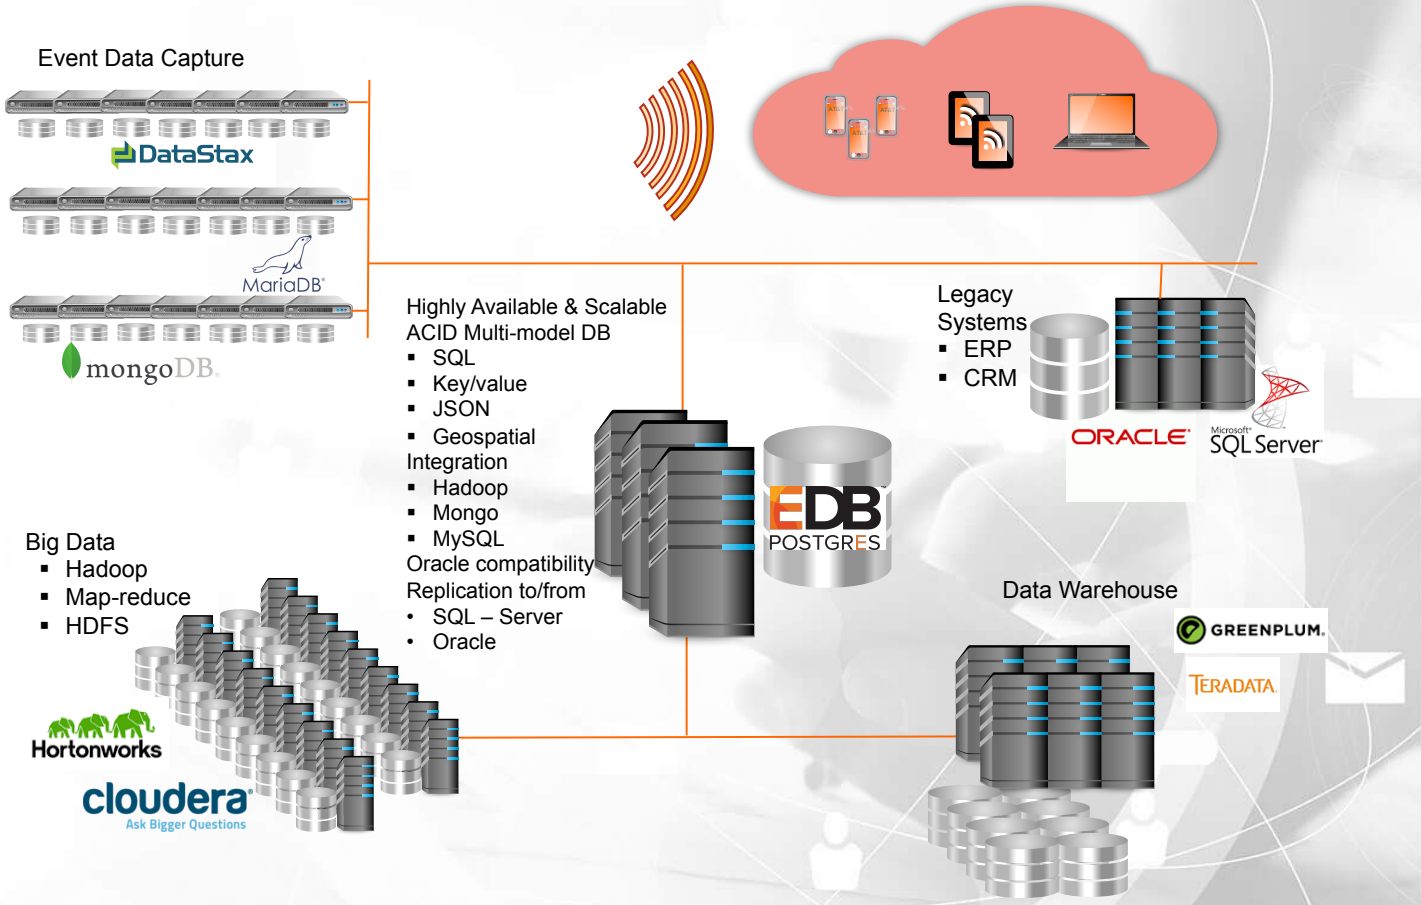

EDB POSTGRES

Focus on Innovation - EDB Postgres Ark – DBaaS Infrastructure

- Flexible DBaaS provisioning for Postgres in multi-platform clouds

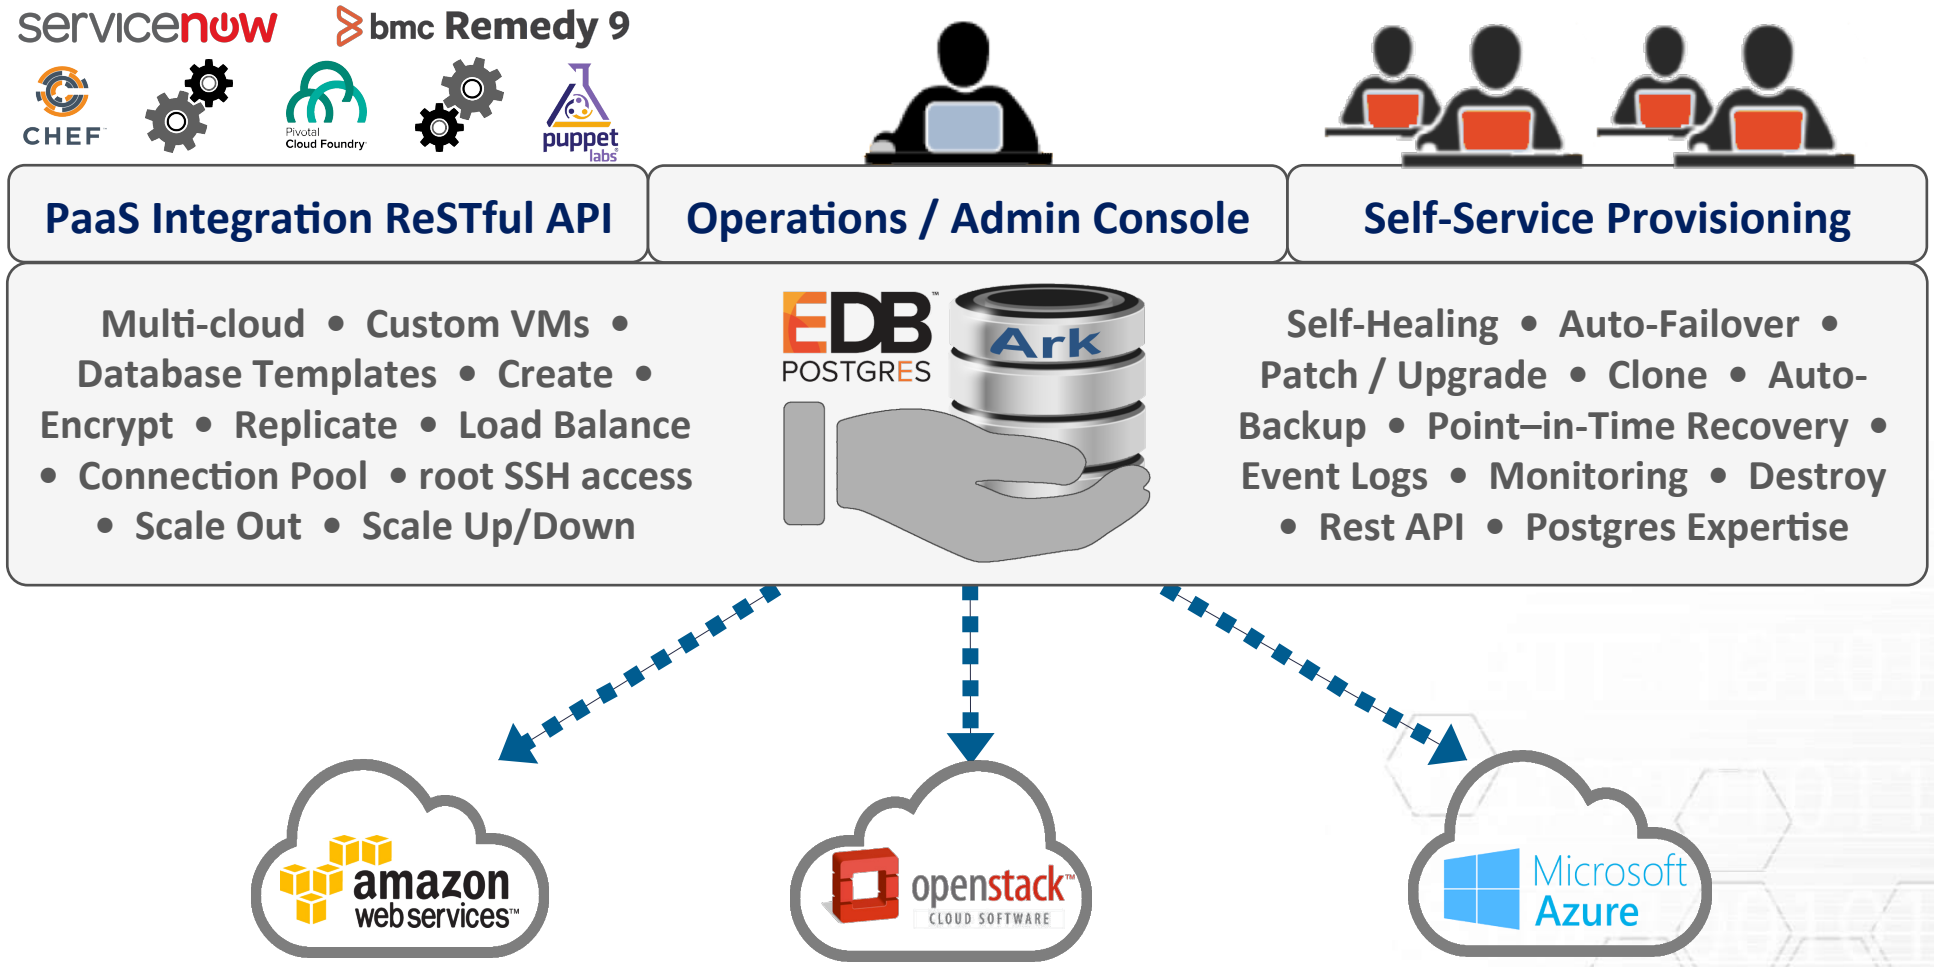

EDB POSTGRES

Focus on Innovation – 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

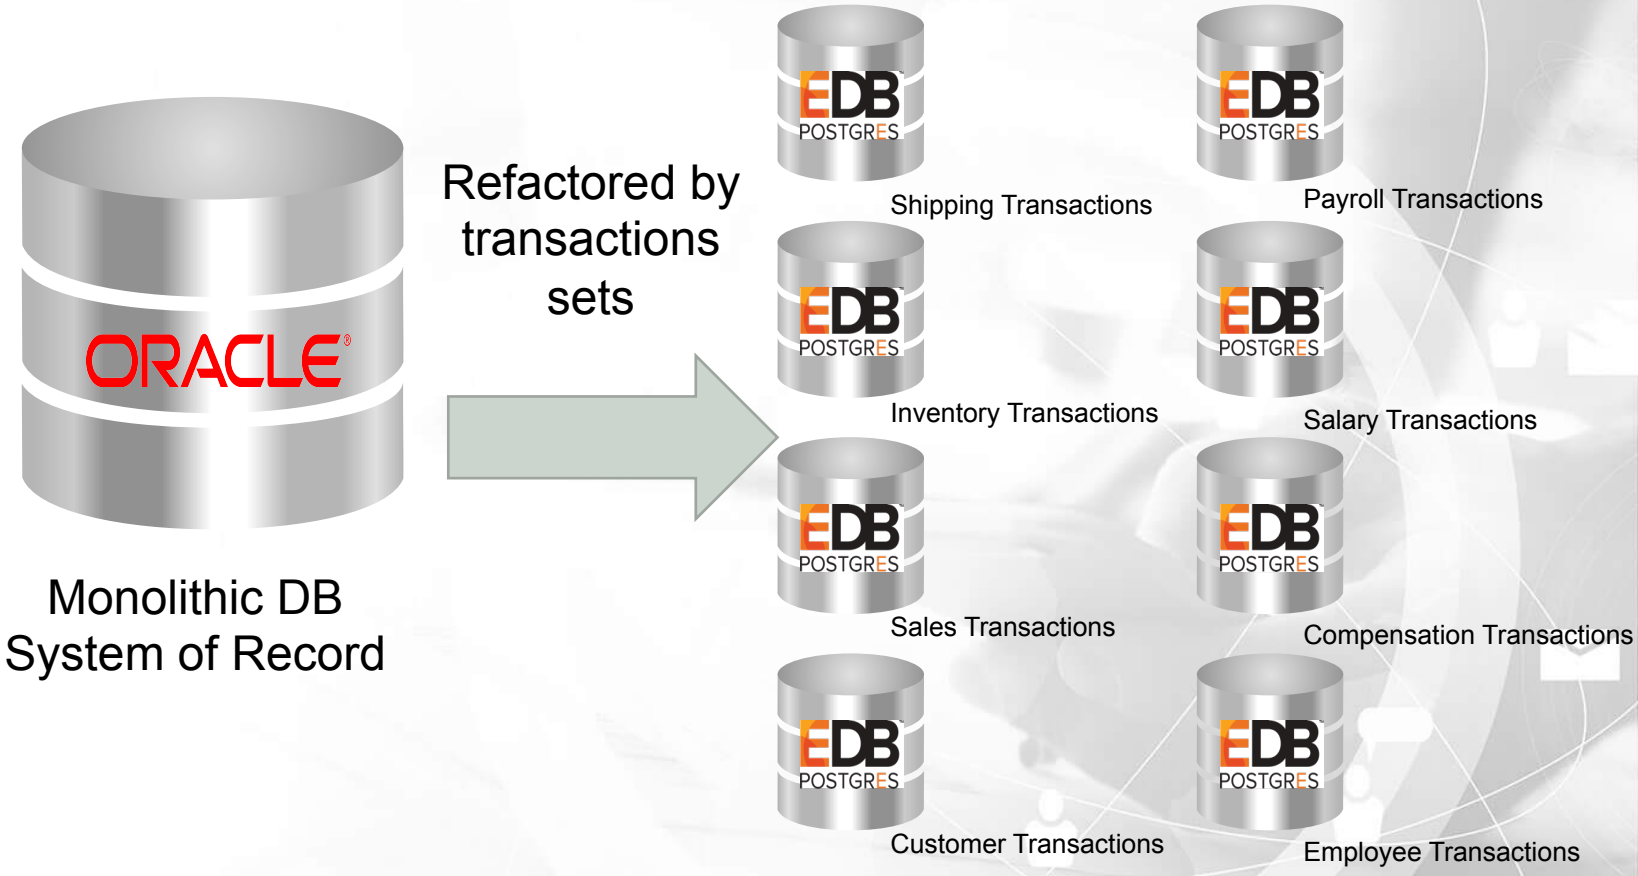

EDB POSTGRES

EnterpriseDB – Postgres Database Containers

Getting the Elephant in the Container...

Toka, a 42 year old [Postgres tatooed] African elephant, is photographed on May 1 2012 at the Metro Toronto Zoo getting accustomed to the container that will be used to send her to another facility [via micro-service].

http://www.huffingtonpost.ca/2012/07/17/toronto-zoo-elephants-california_n_1681380.html

Turns out, Elephants love containers!

Postgres in a container, done...

```
FROM centos:latest

MAINTAINER <email address>

RUN rpm -Uvh http://yum.enterprisedb.com/edbrepos/edb-repo-9.6-4.noarch.rpm \
  && export YUM_USER=<yum user> \
  && export YUM_PASSWORD=<yum password> \
  && sed -i "s/<username>:<password>/$YUM_USER:$YUM_PASSWORD/g" /etc/yum.repos.d/edb.repo \
  && sed -i "\/ppas95/,/gpgcheck/ s/enabled=0/enabled=1/" /etc/yum.repos.d/edb.repo \
  && yum -y install ppas95-server

USER enterprisedb

RUN echo "enterprisedb" > /tmp/password \
   && /usr/ppas-9.5/bin/initdb -D /var/lib/ppas/9.5/data --pwfile=/tmp/password \
   && echo "host   all   all   0.0.0.0/0   md5" >> /var/lib/ppas/9.5/data/pg_hba.conf

CMD ["/usr/ppas-9.5/bin/edb-postgres", "-D", "/var/lib/ppas/9.5/data", "-h", "*"]
```

Ok, that’s cute, but what about...

- Backup
- Recovery
- Replication
- Load Balancing
- Failover
- ...

Let’s start with some fundamentals...

What is a database?

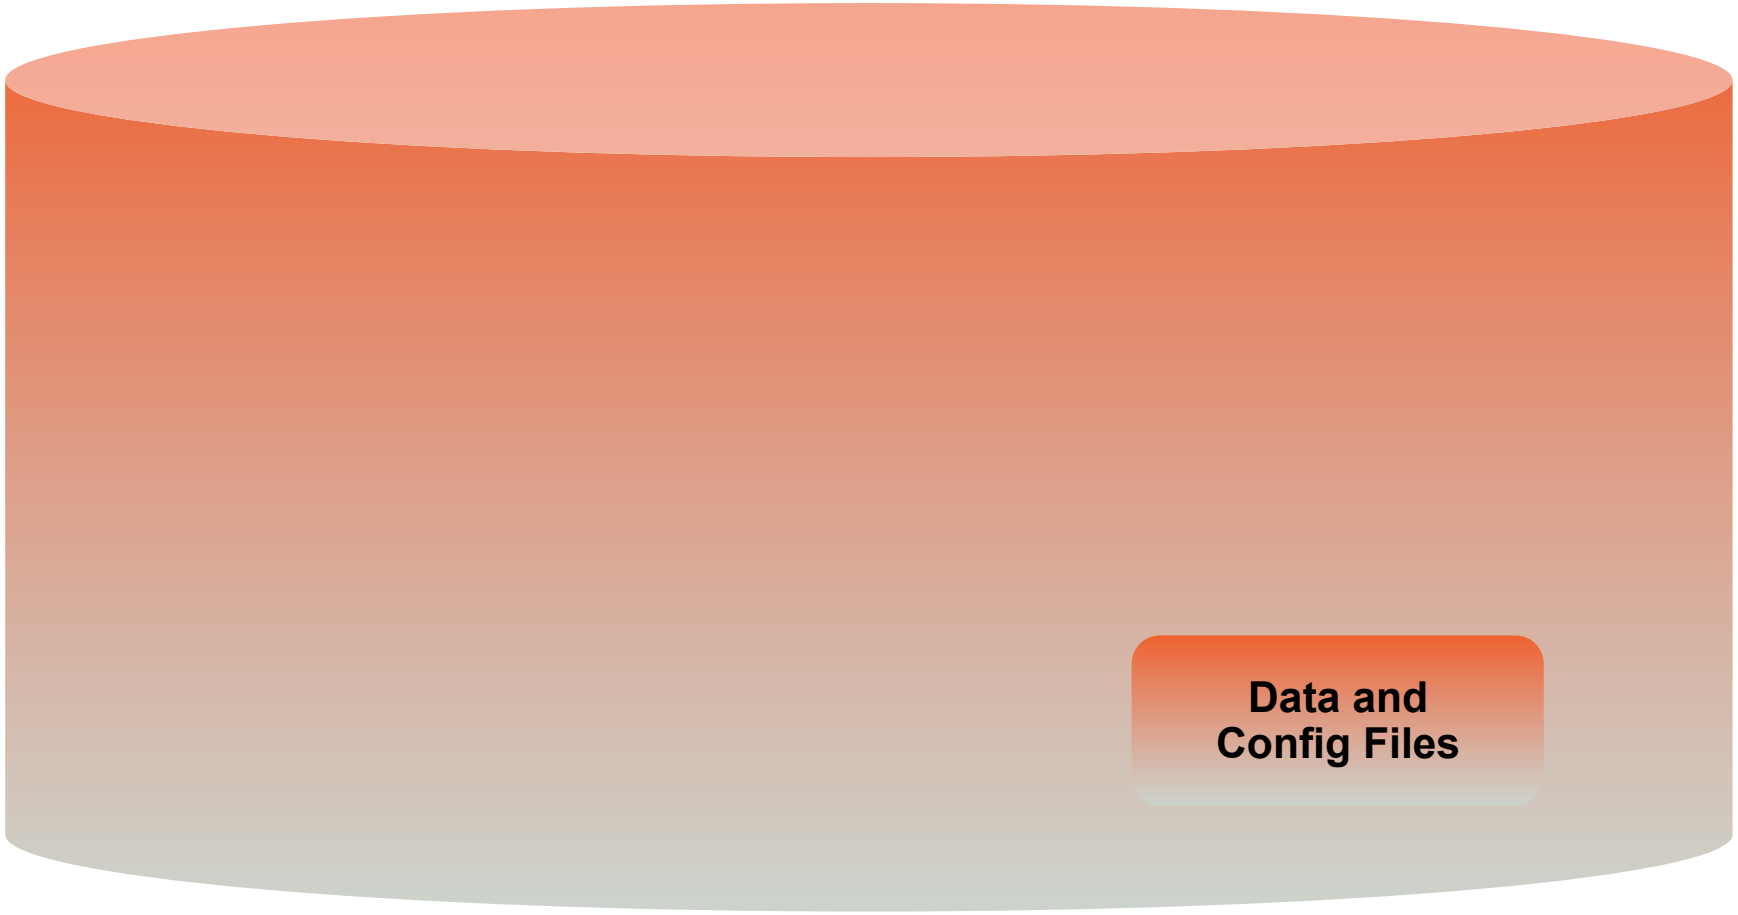

What is a database?

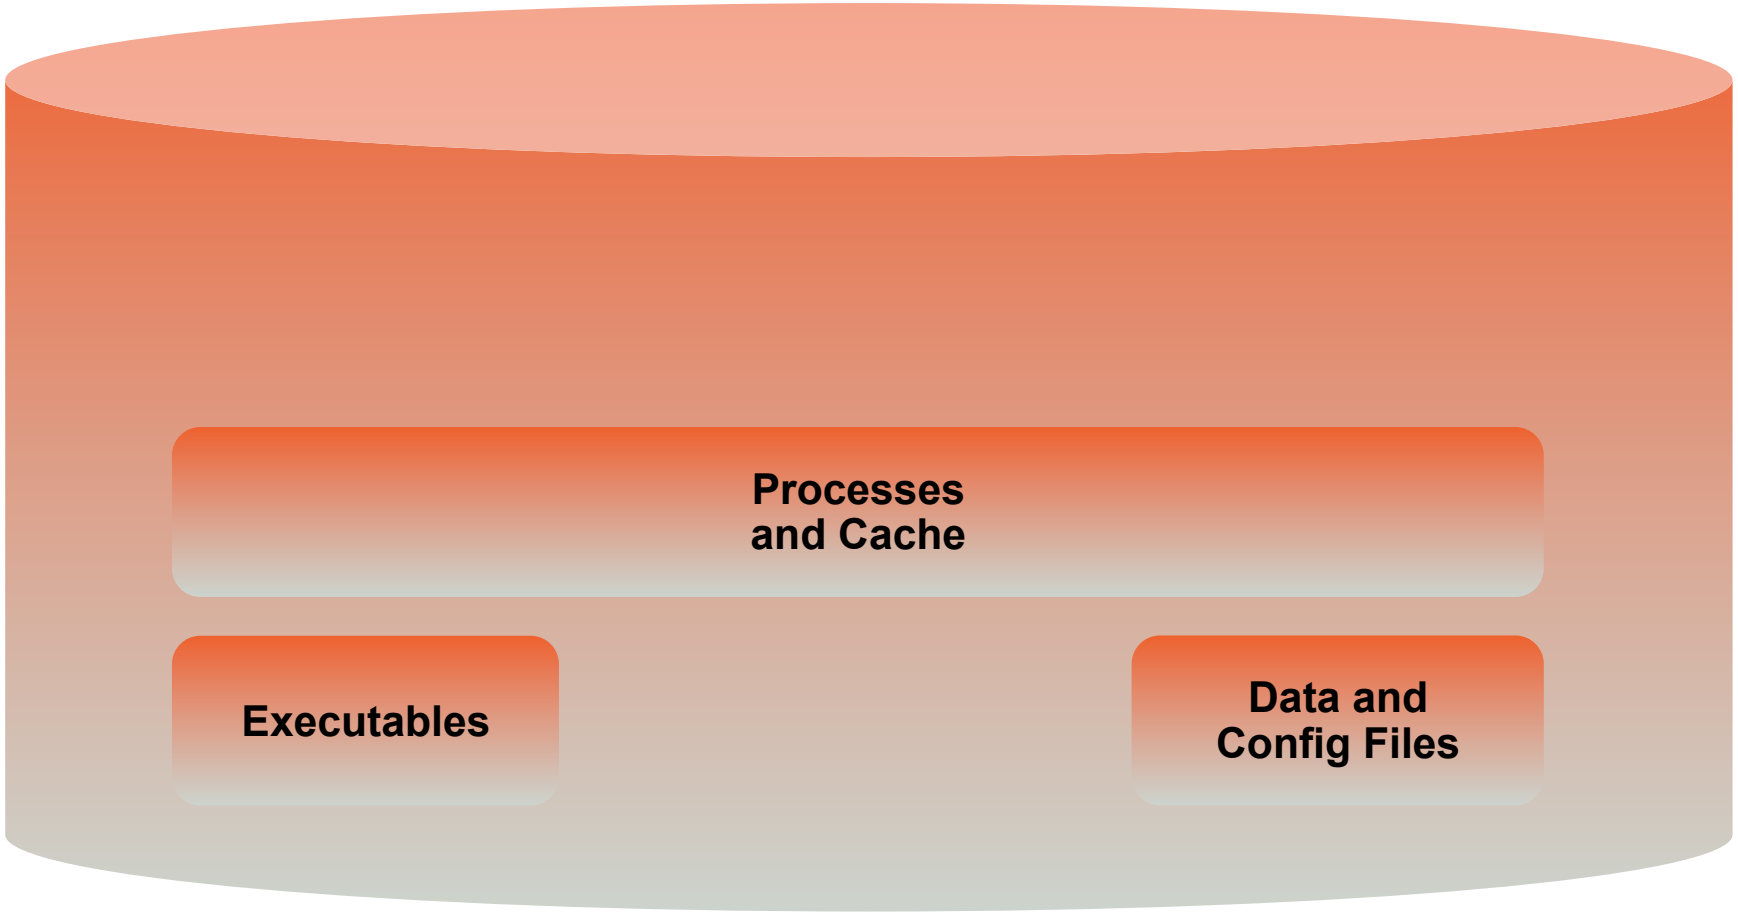

What is a database?

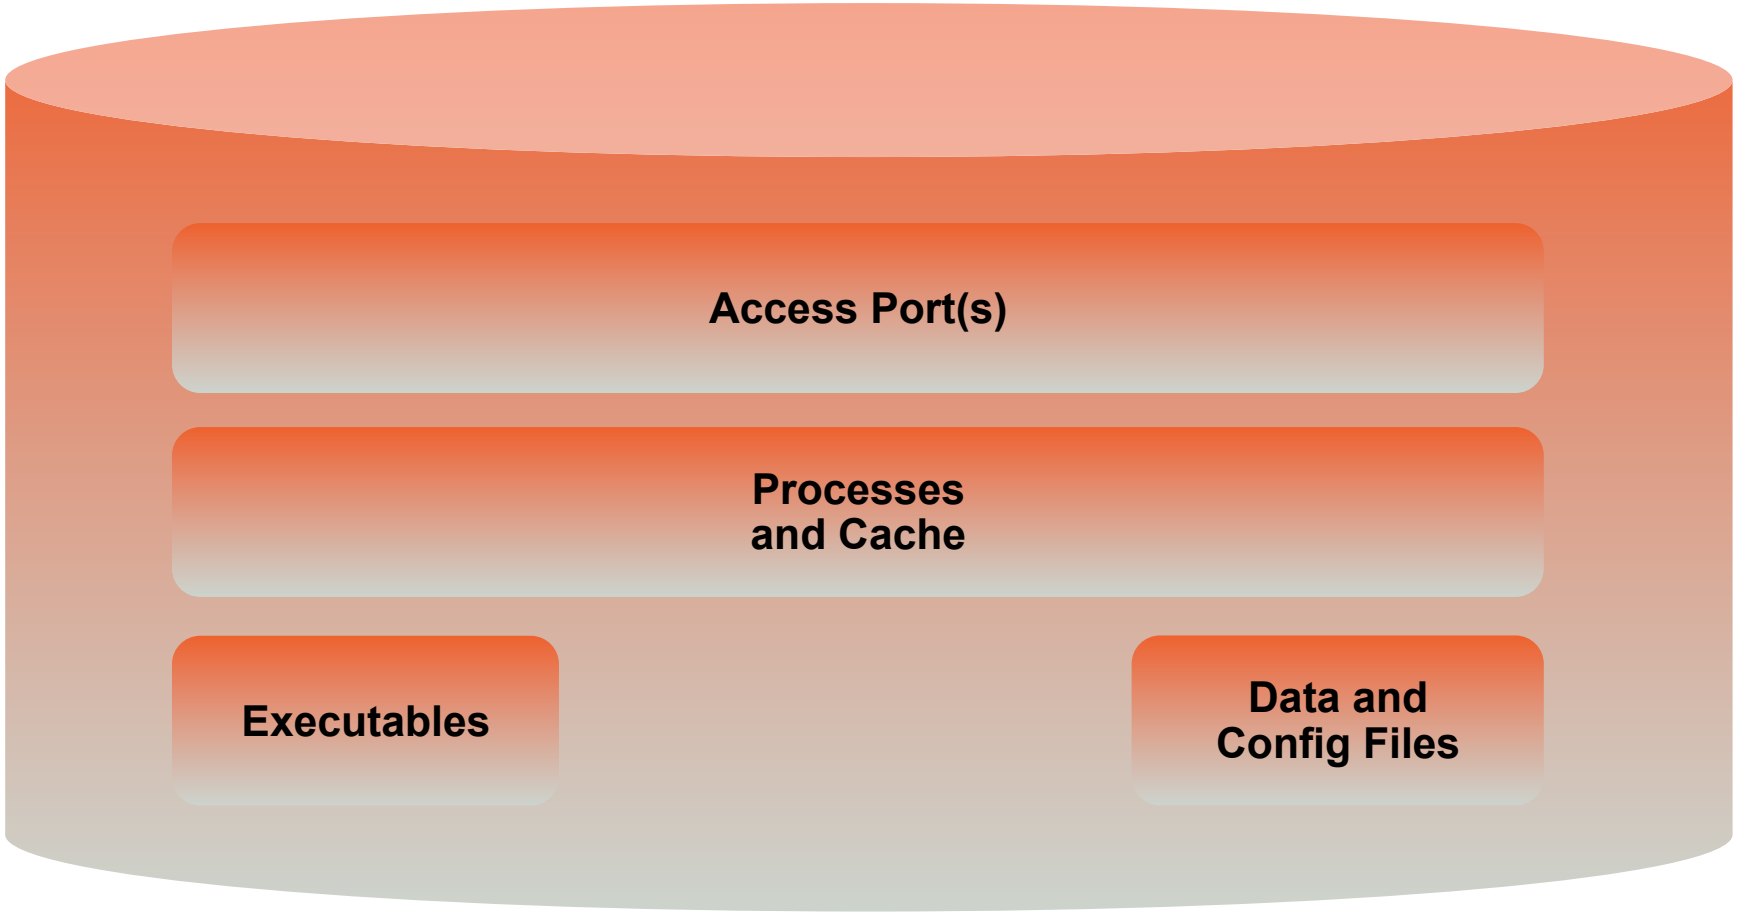

What is a database?

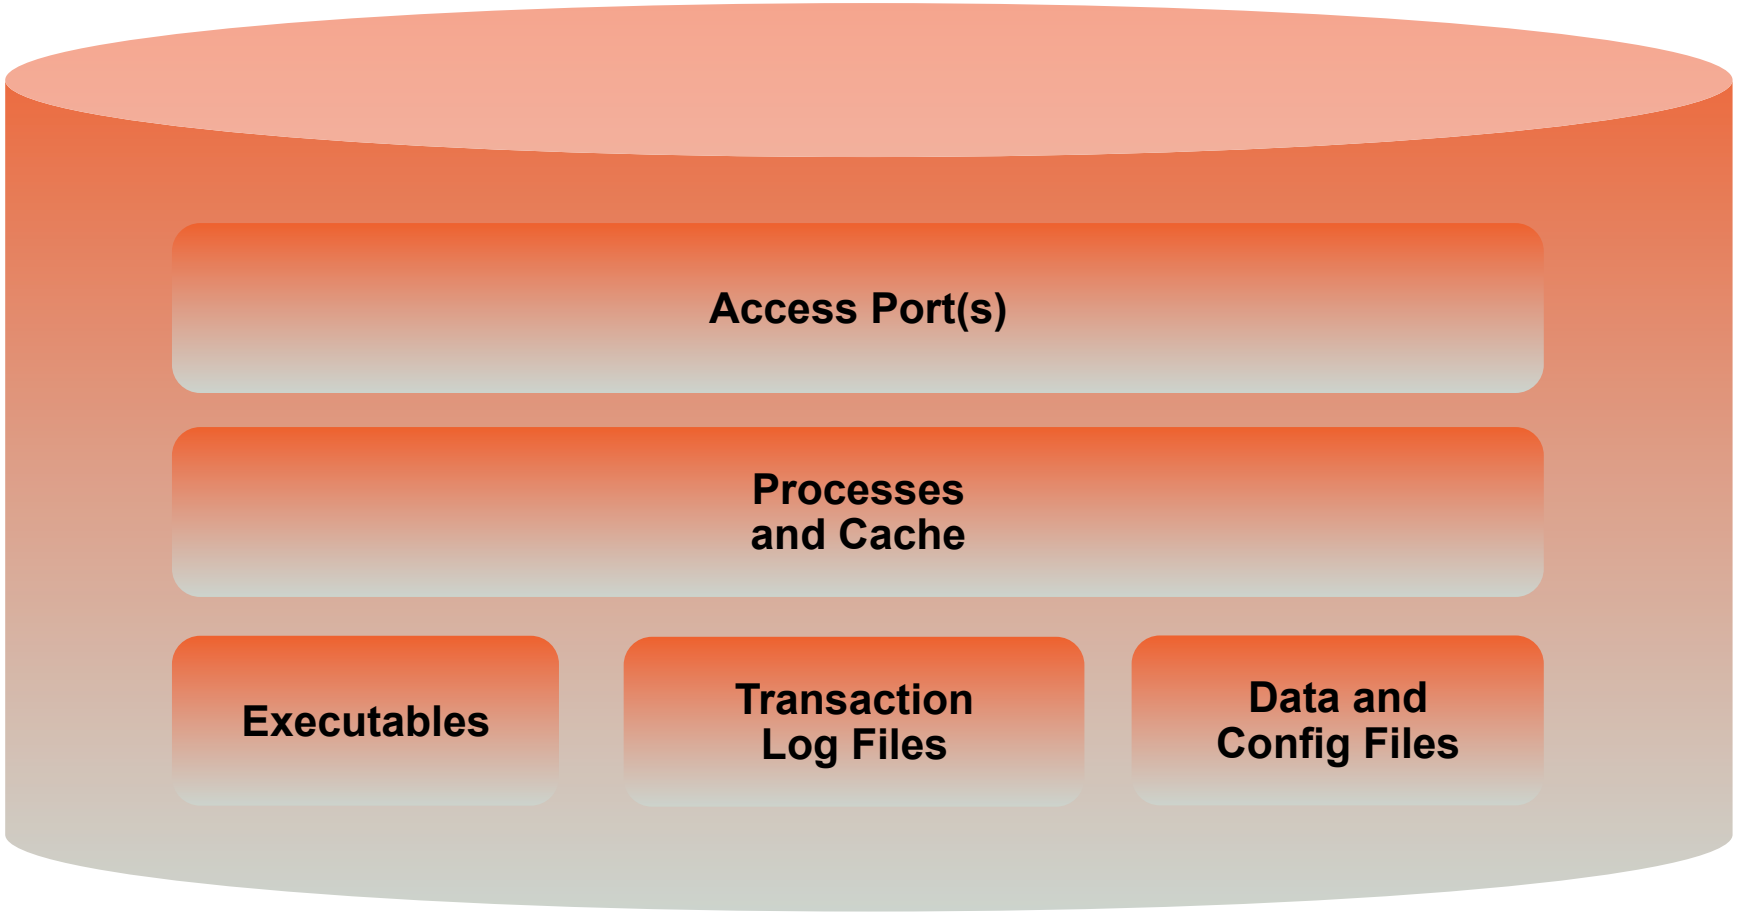

What is an HA database setup?

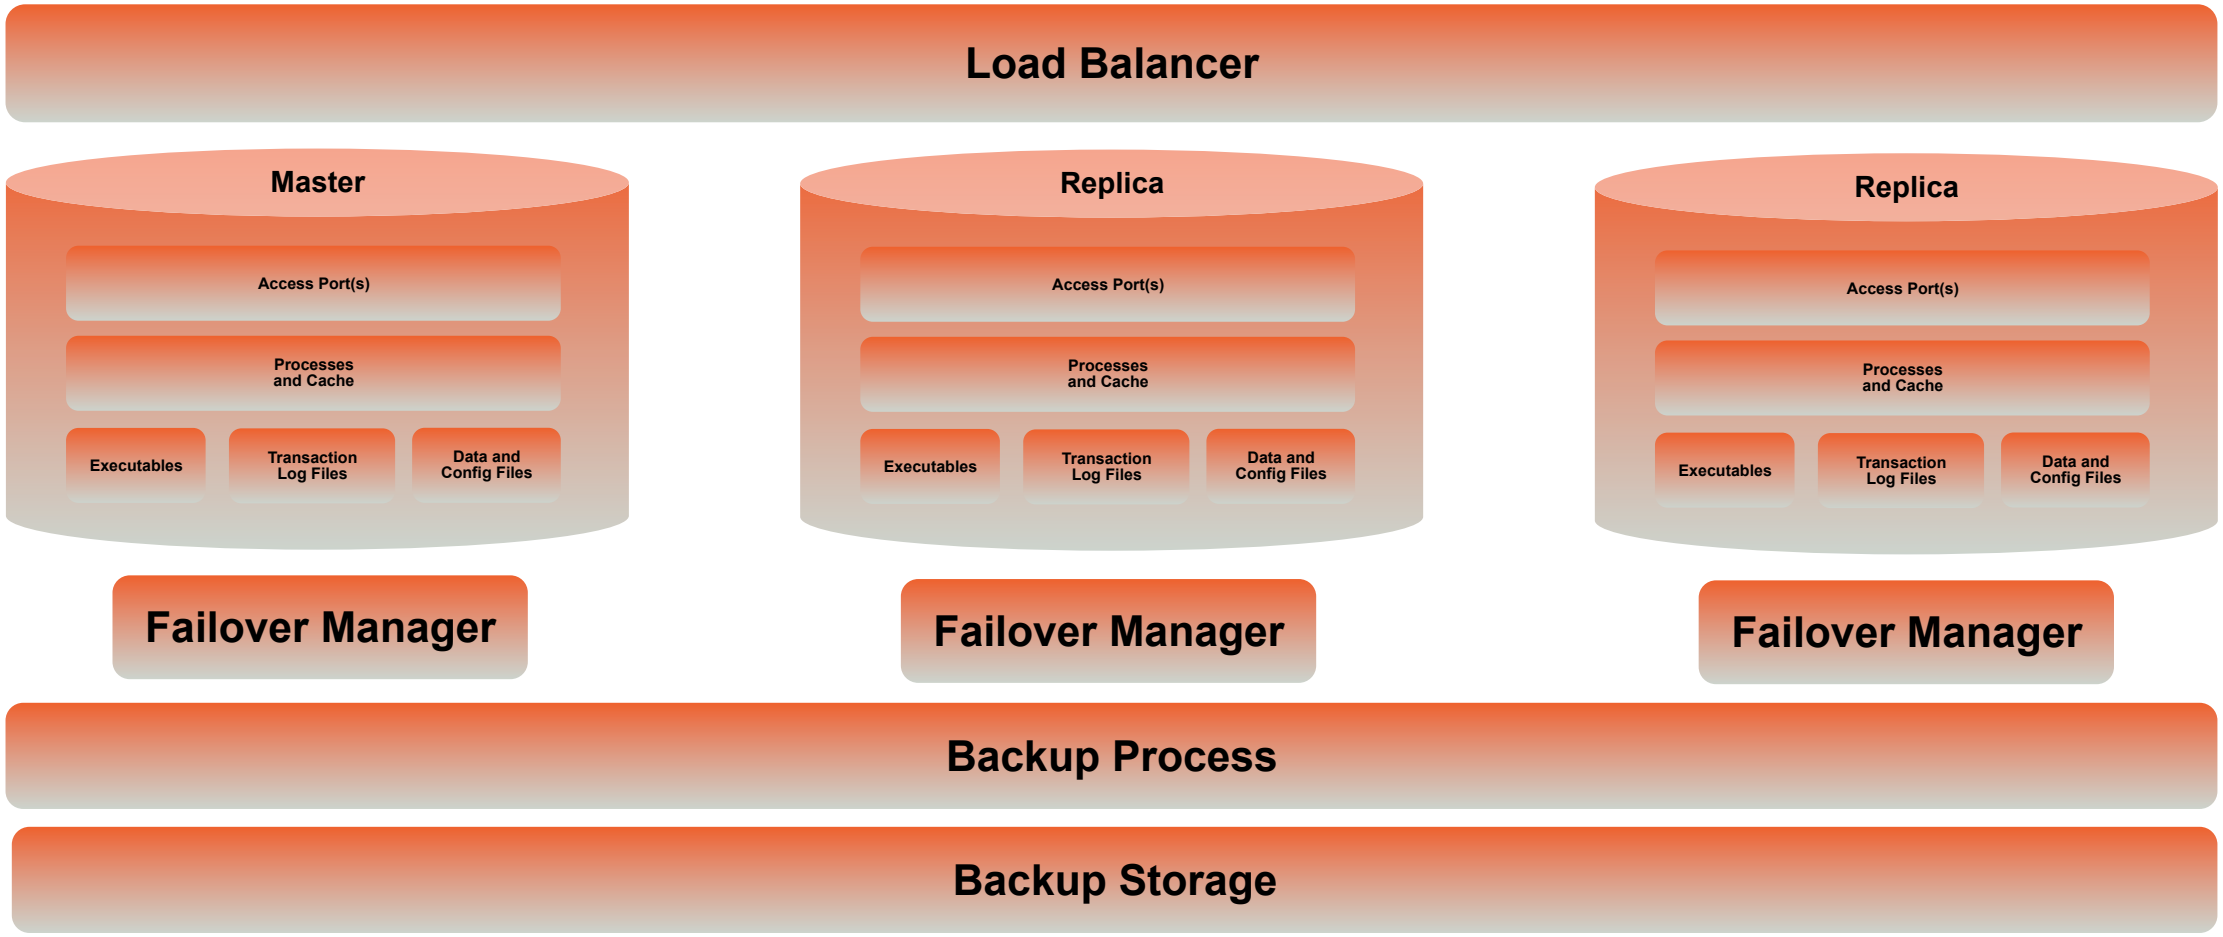

So, what does one containerize?

- Different strokes for different folks...

The Mother Container...

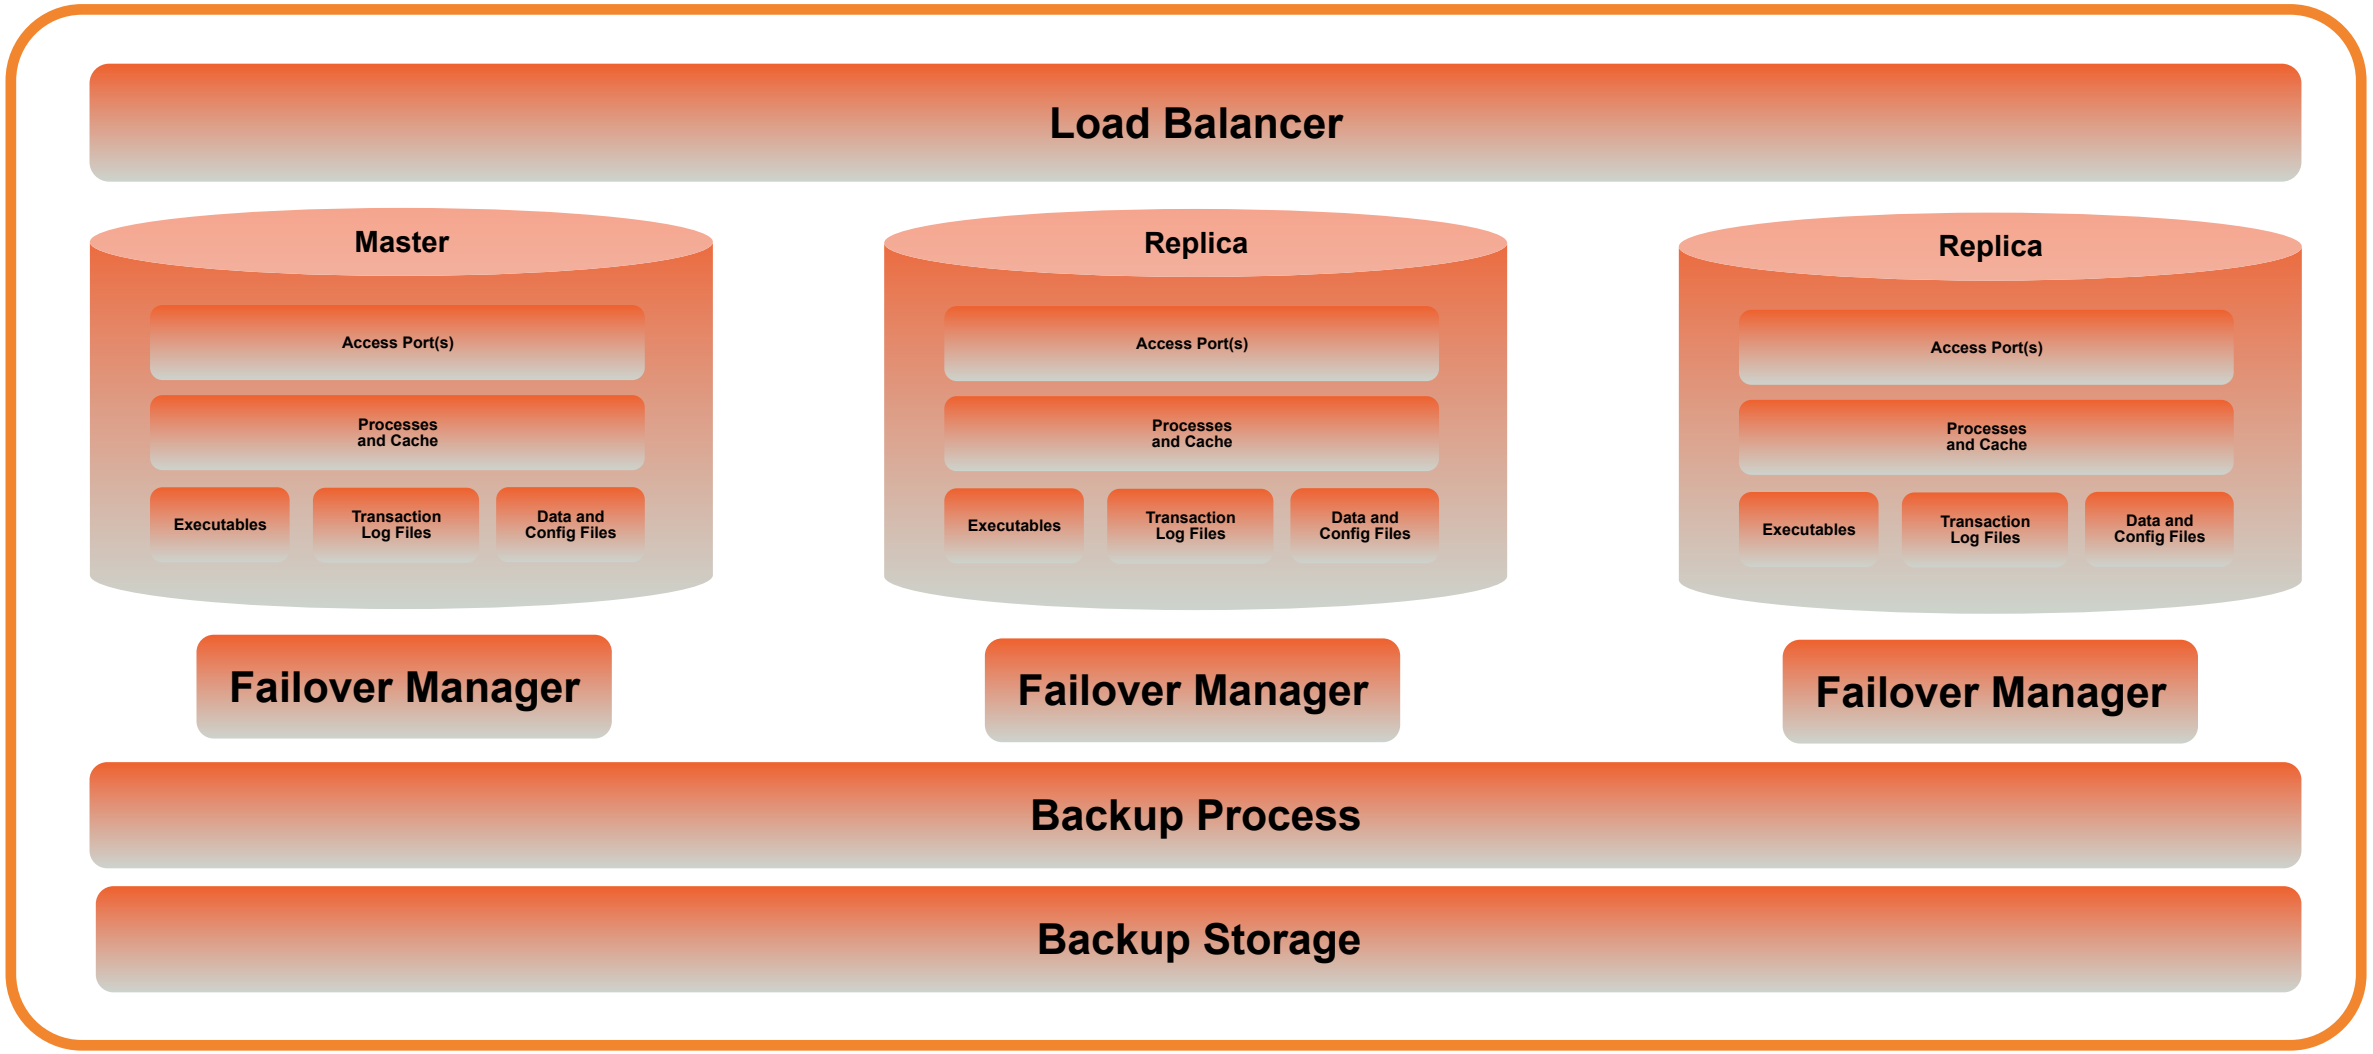

The Academic Container(s)...

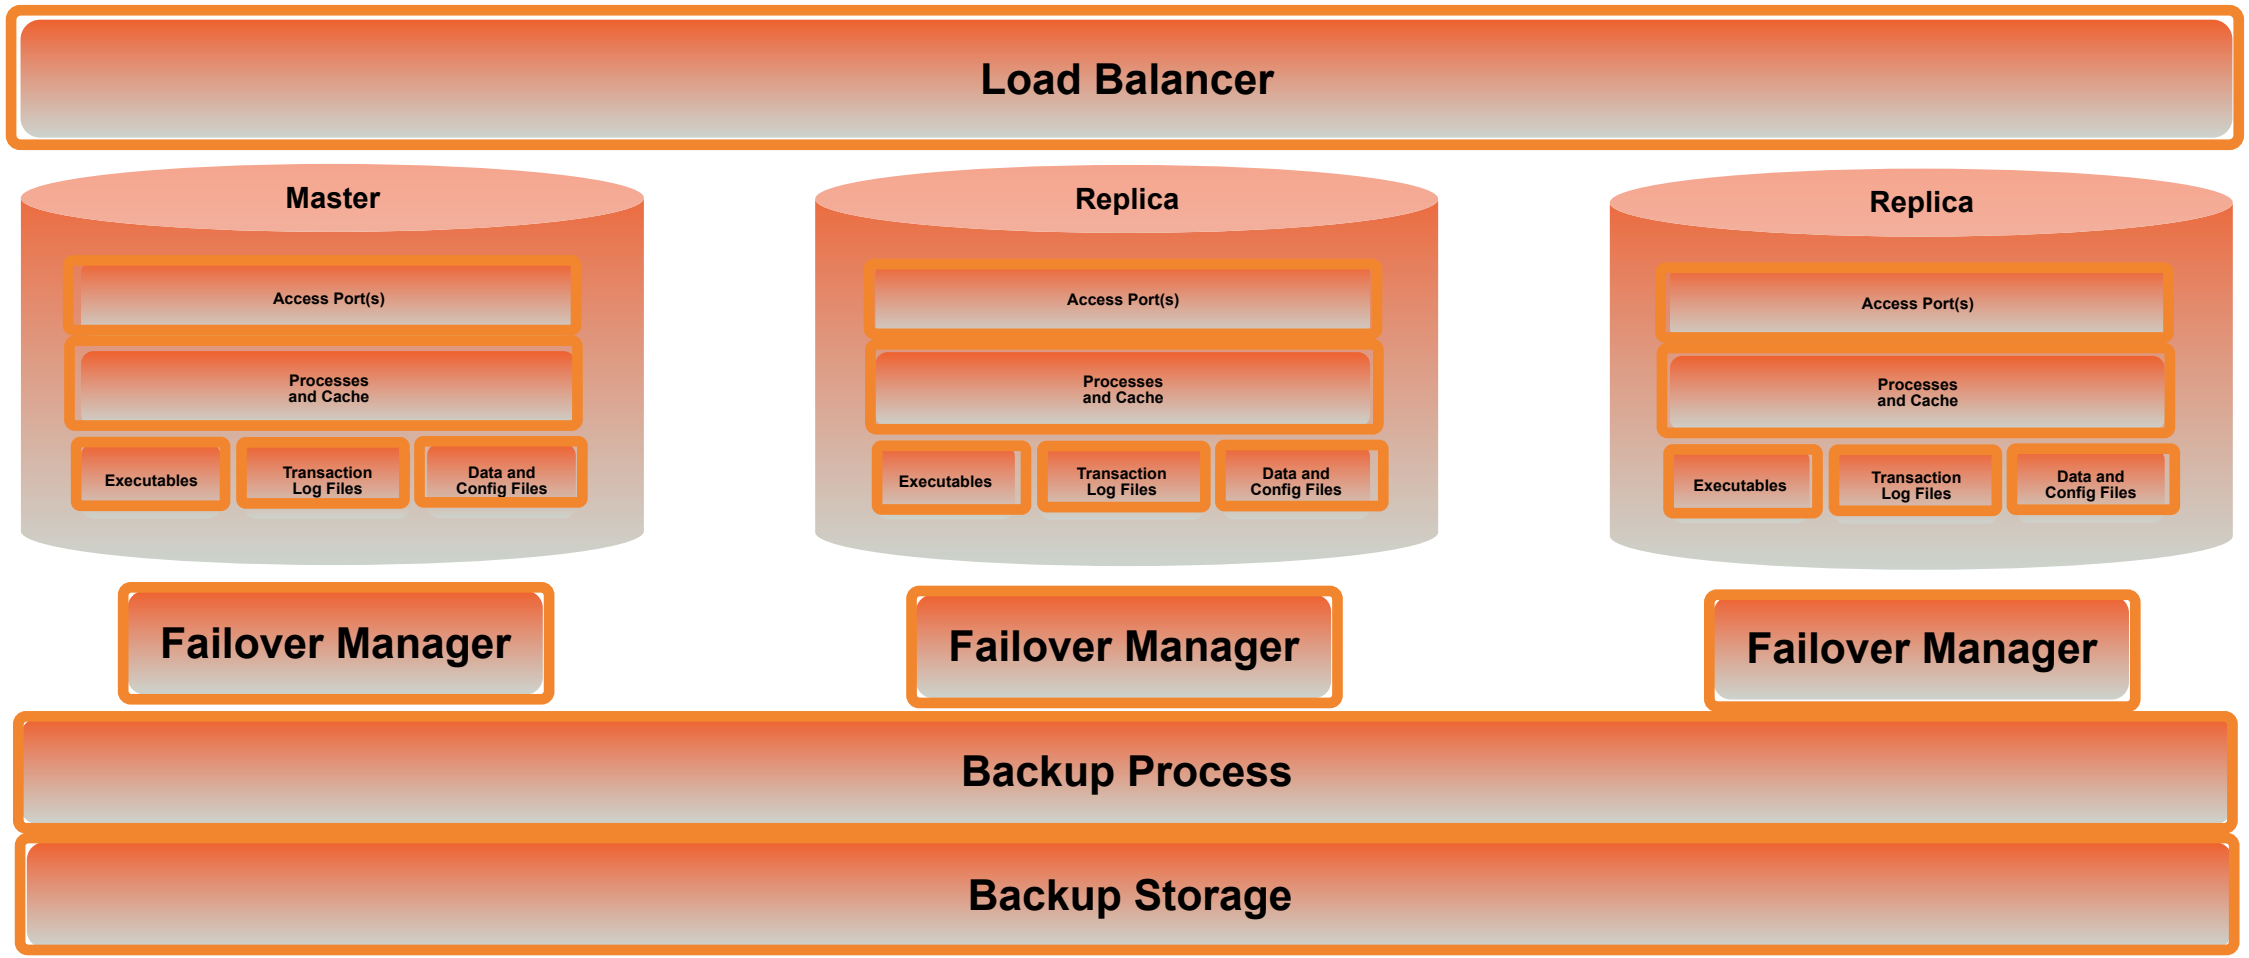

Now, let's give this some thought...

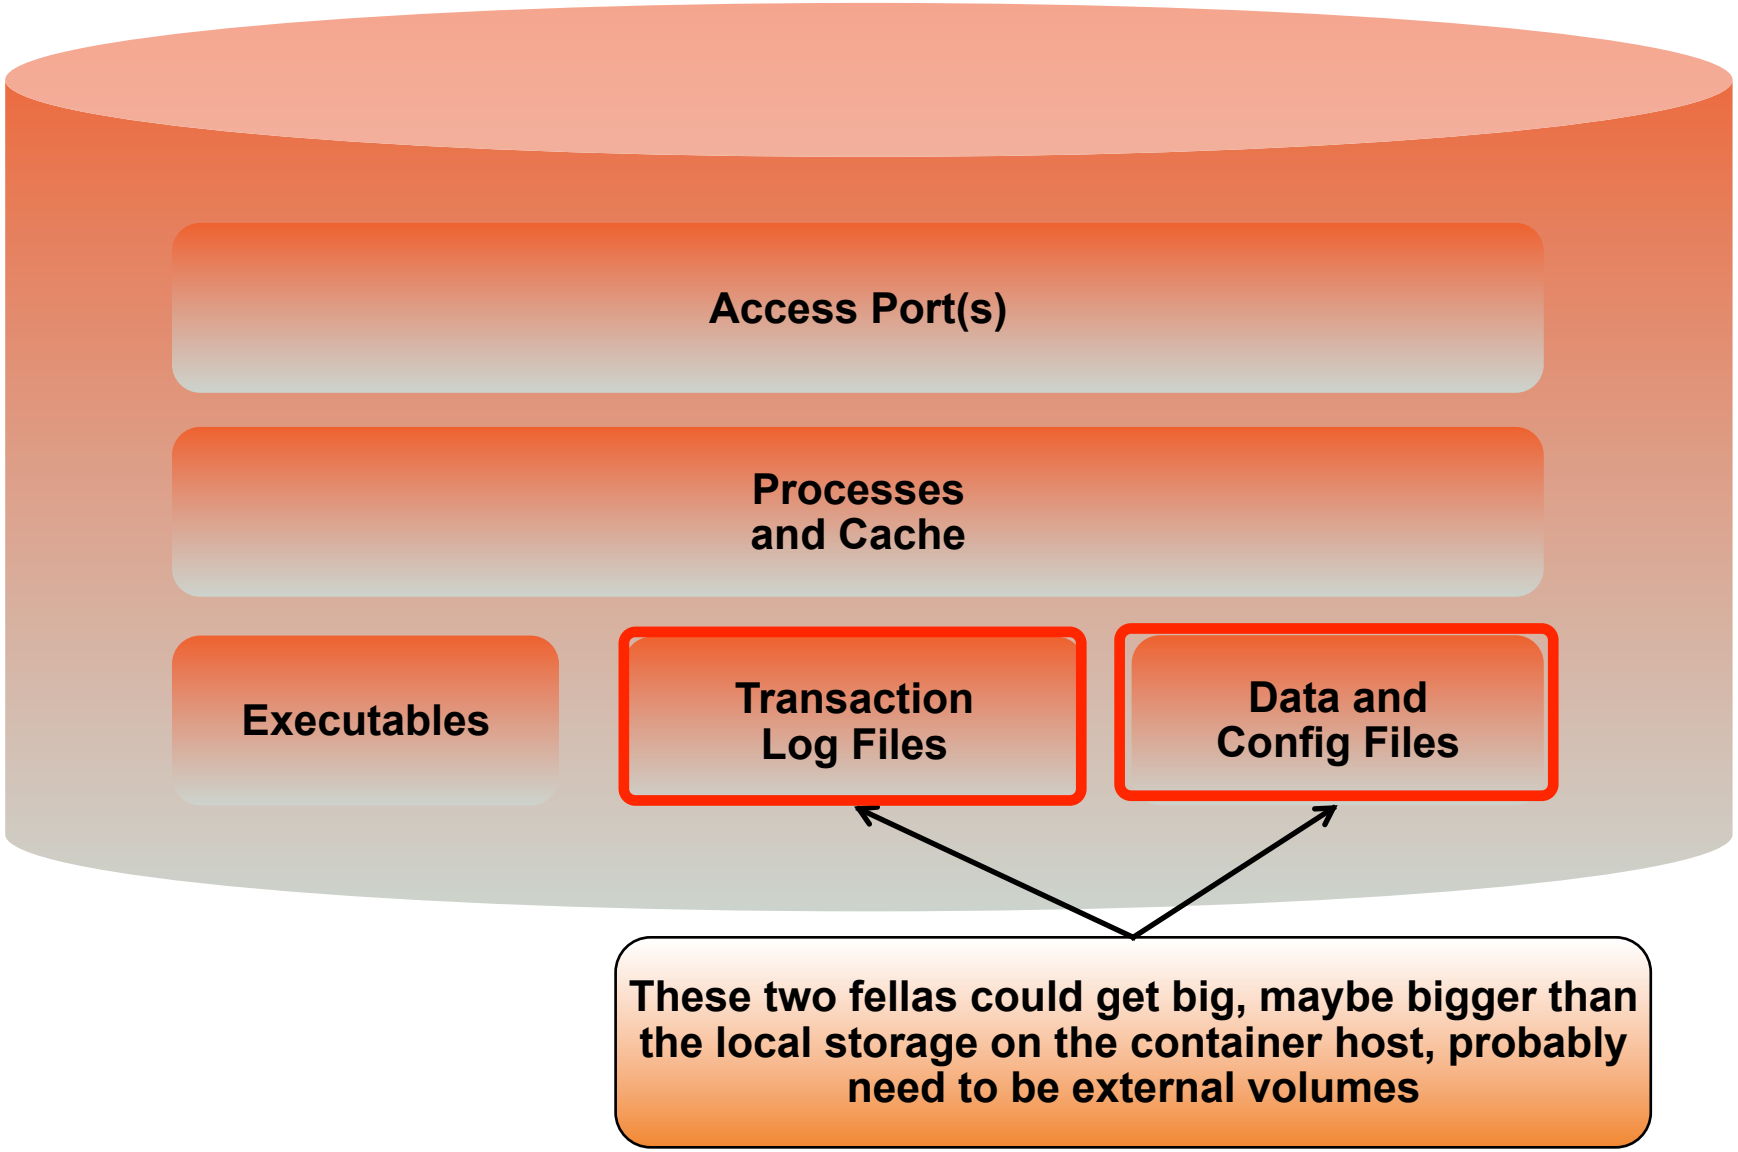

Pragmatic attempt #1

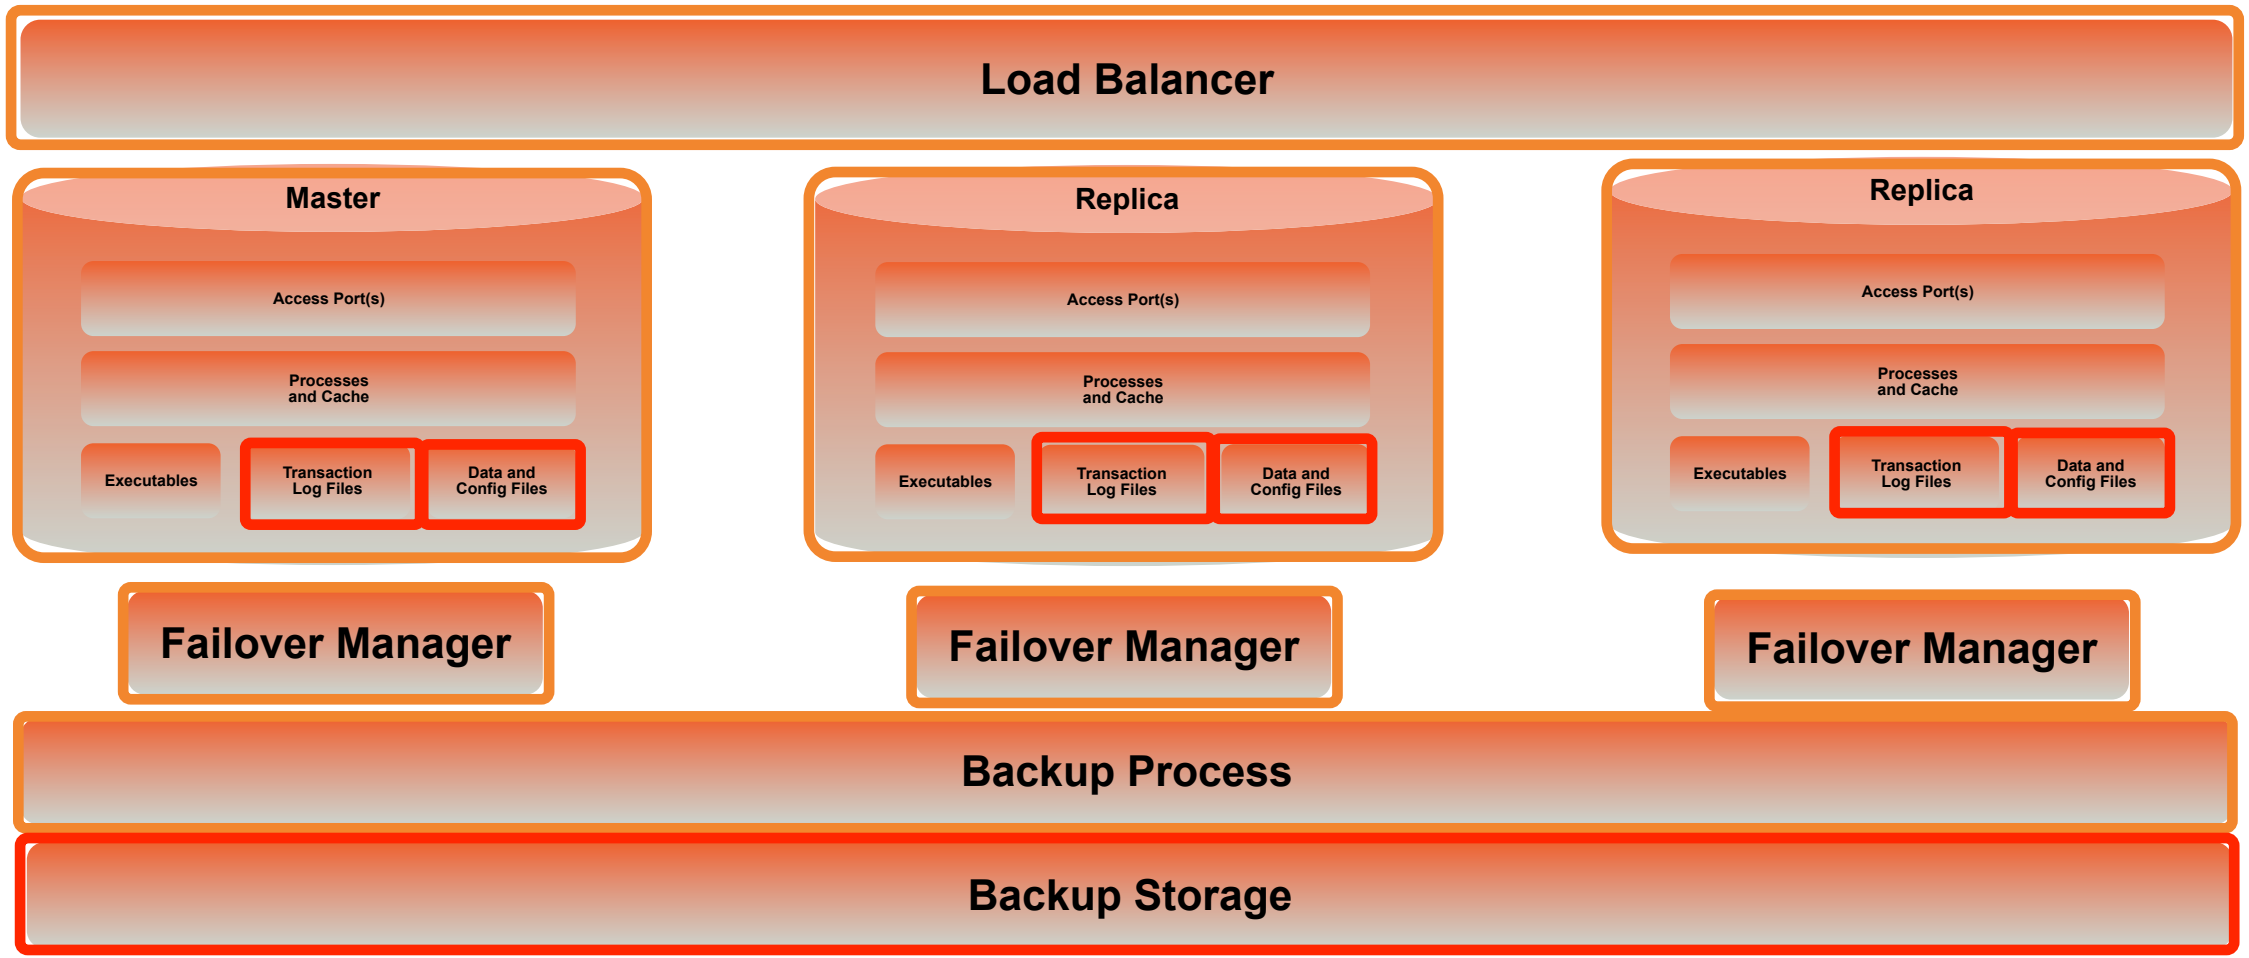

Not bad, but...

- The pesky master/replica issue
  - One of your DB containers is a master, the others are replicas, and it's pretty important to know the difference
- The attached at the hip issue
  - Failover managers probably need to be on each node one to one to handle promotions and reconfiguration of replicas to point to new masters
  - Kubernetes/OpenShift Pods can help with this, but that limits you to that environment, and you might have to use some admin APIs that get frowned on
- The single load balancer issue
  - Now you've got to replicate that as well and handle failover, plus updating all of them when masters/replicas change
- The container explosion issue
  - A reasonably HA environment probably has 10+ containers, starting to get complex and possibly fragile

What if we made “smart” containers?

- You mean like the mother container?
- No, more like a mother who married an academic...
- What if...
  - Each container could “pretend” it was a master
  - Each container could quorum with the others to handle failover
  - Each container could work as it’s own atomic unit
- Let’s see what that might look like...

The “Smart” Container...

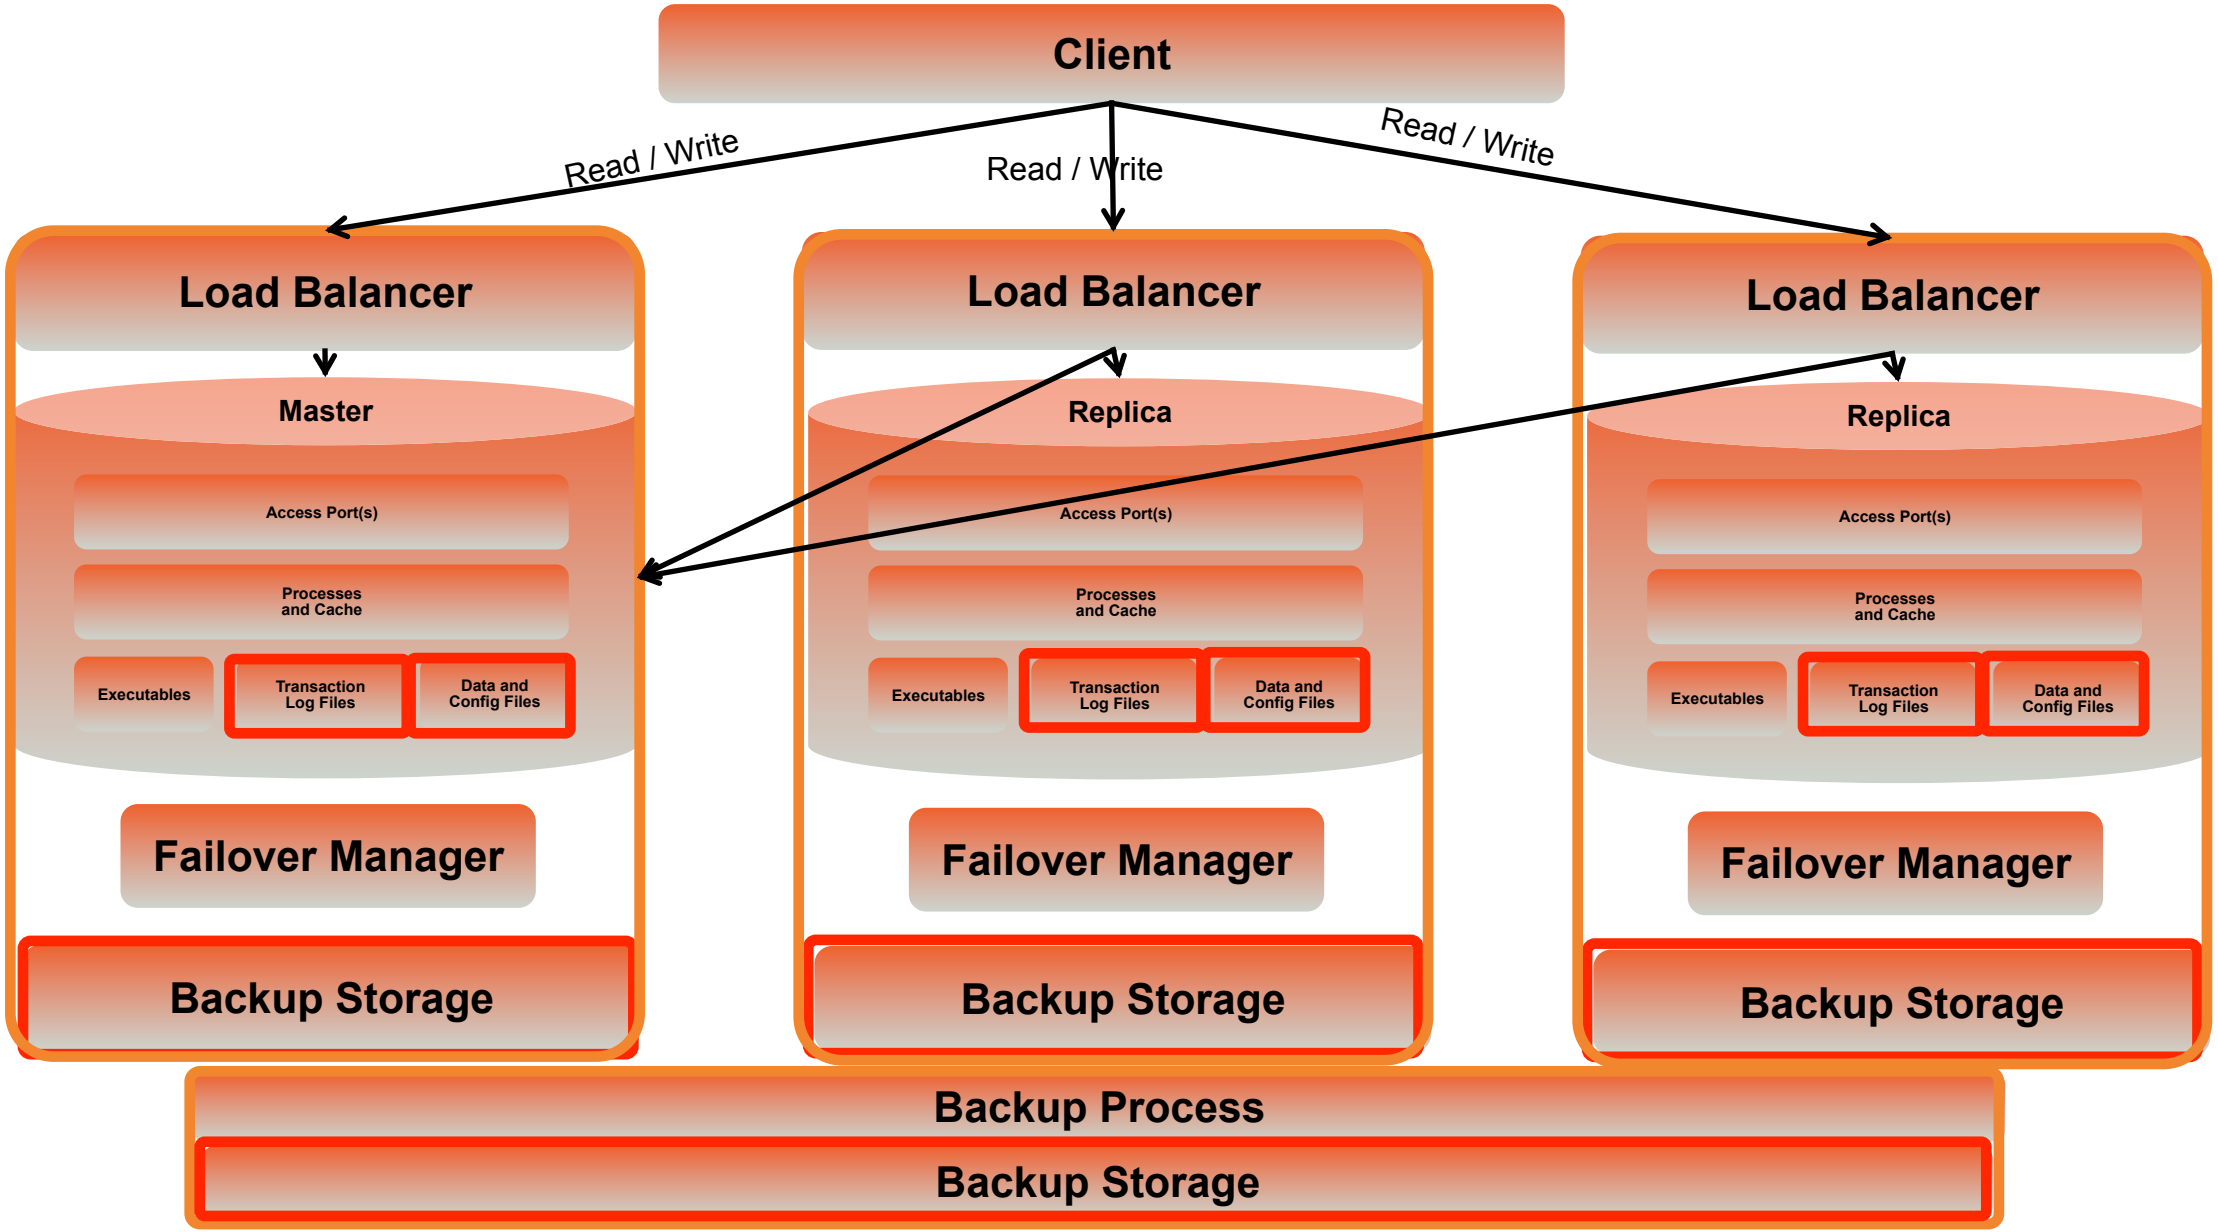

Single Master

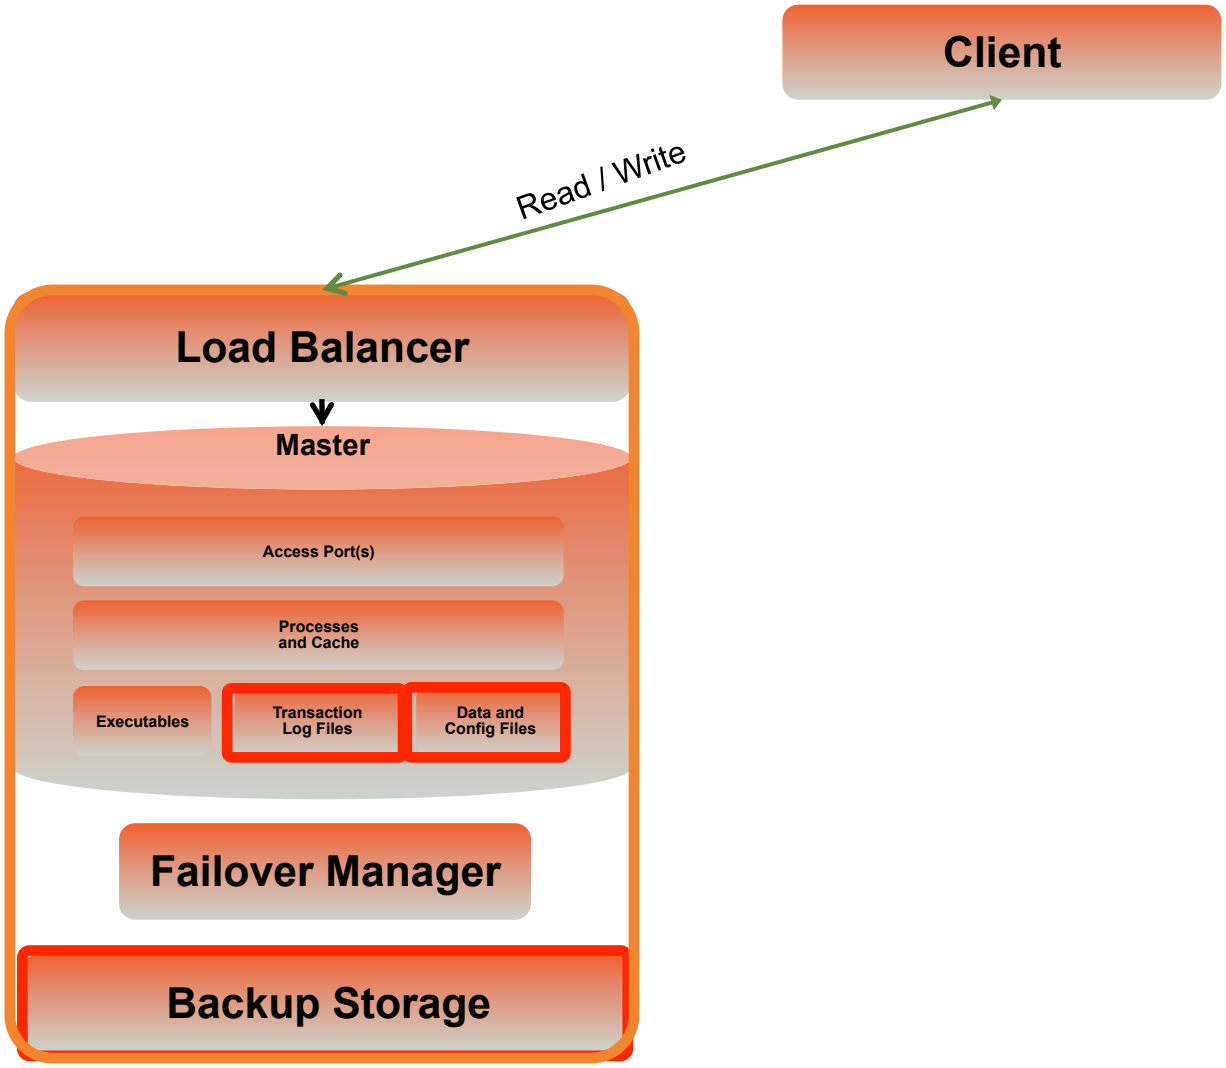

Disaster Recovery

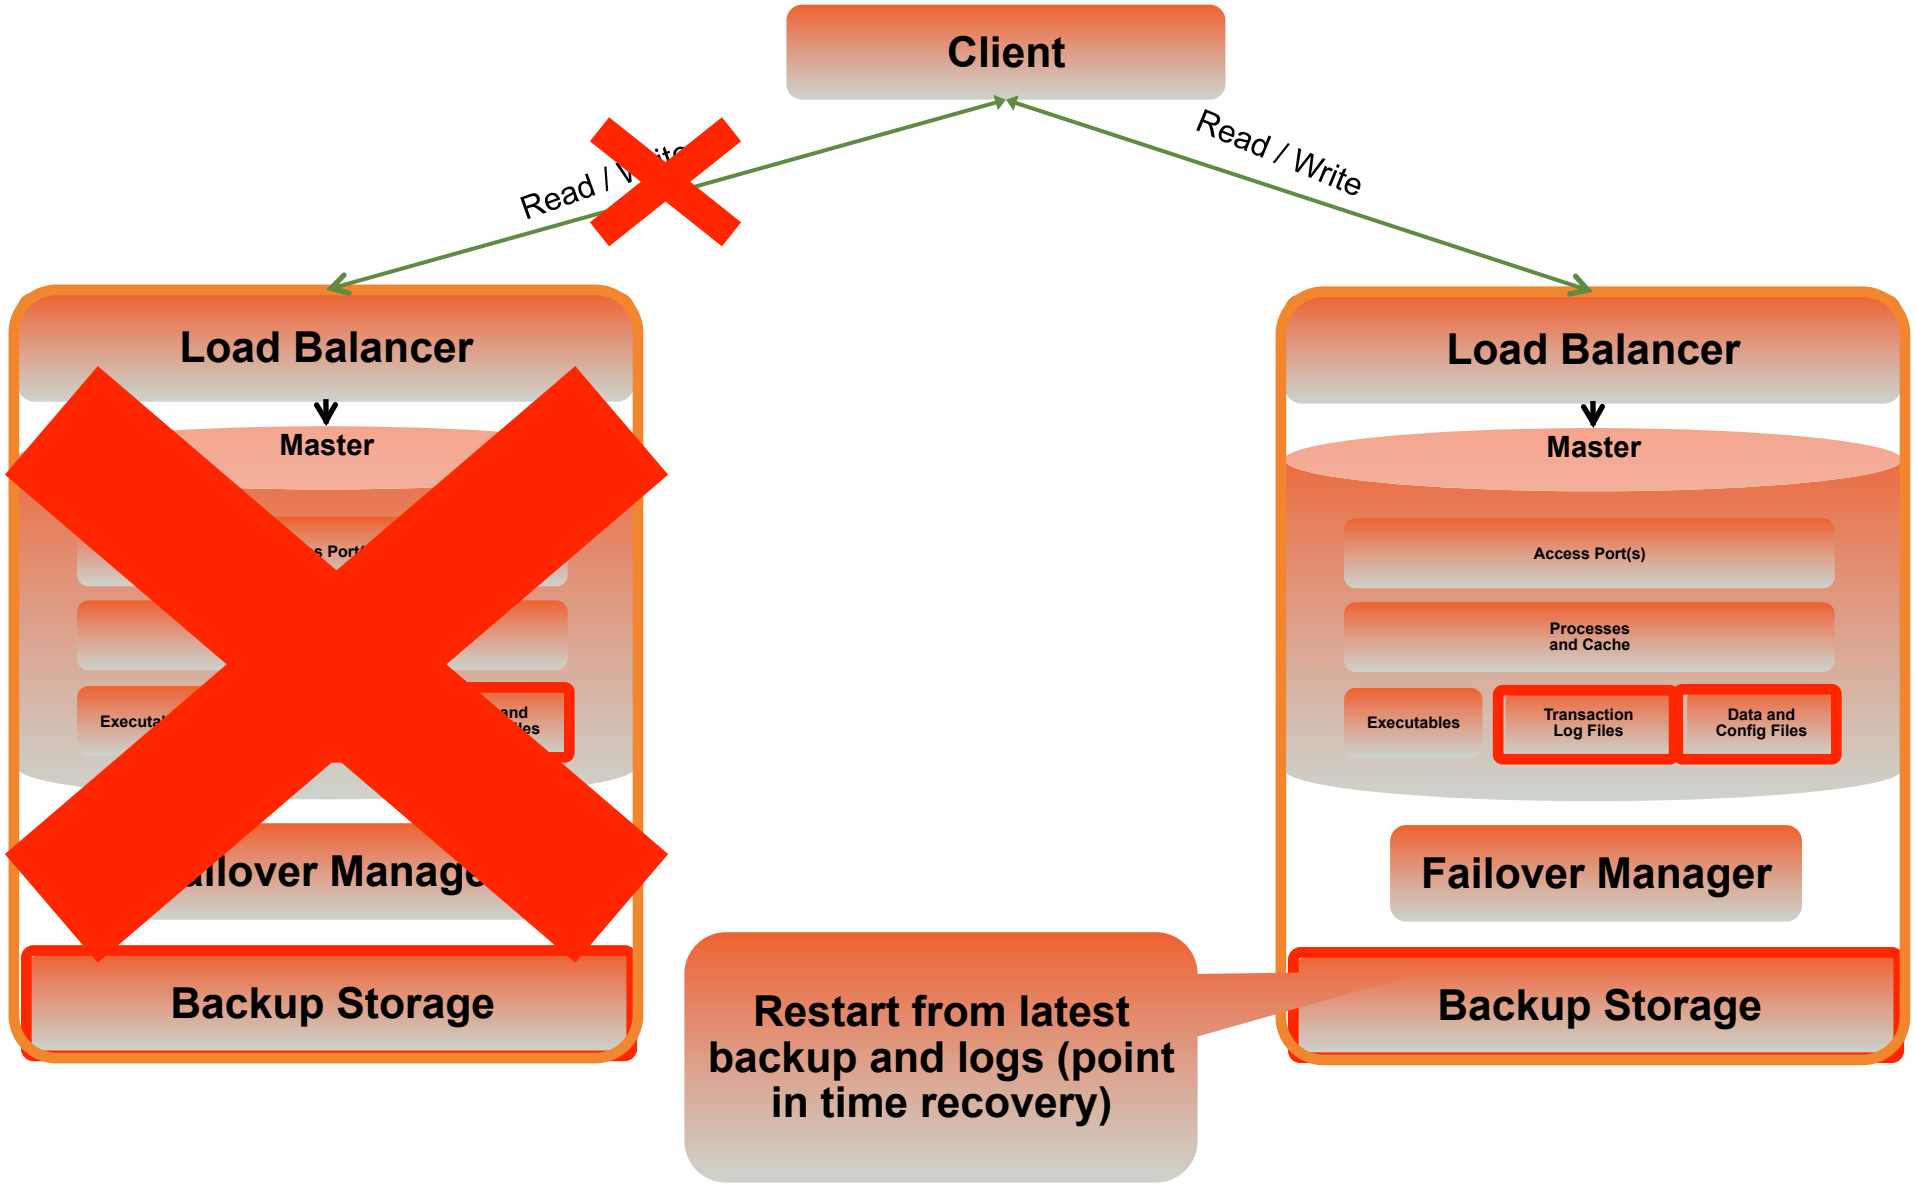

Two “Masters” (scale up)

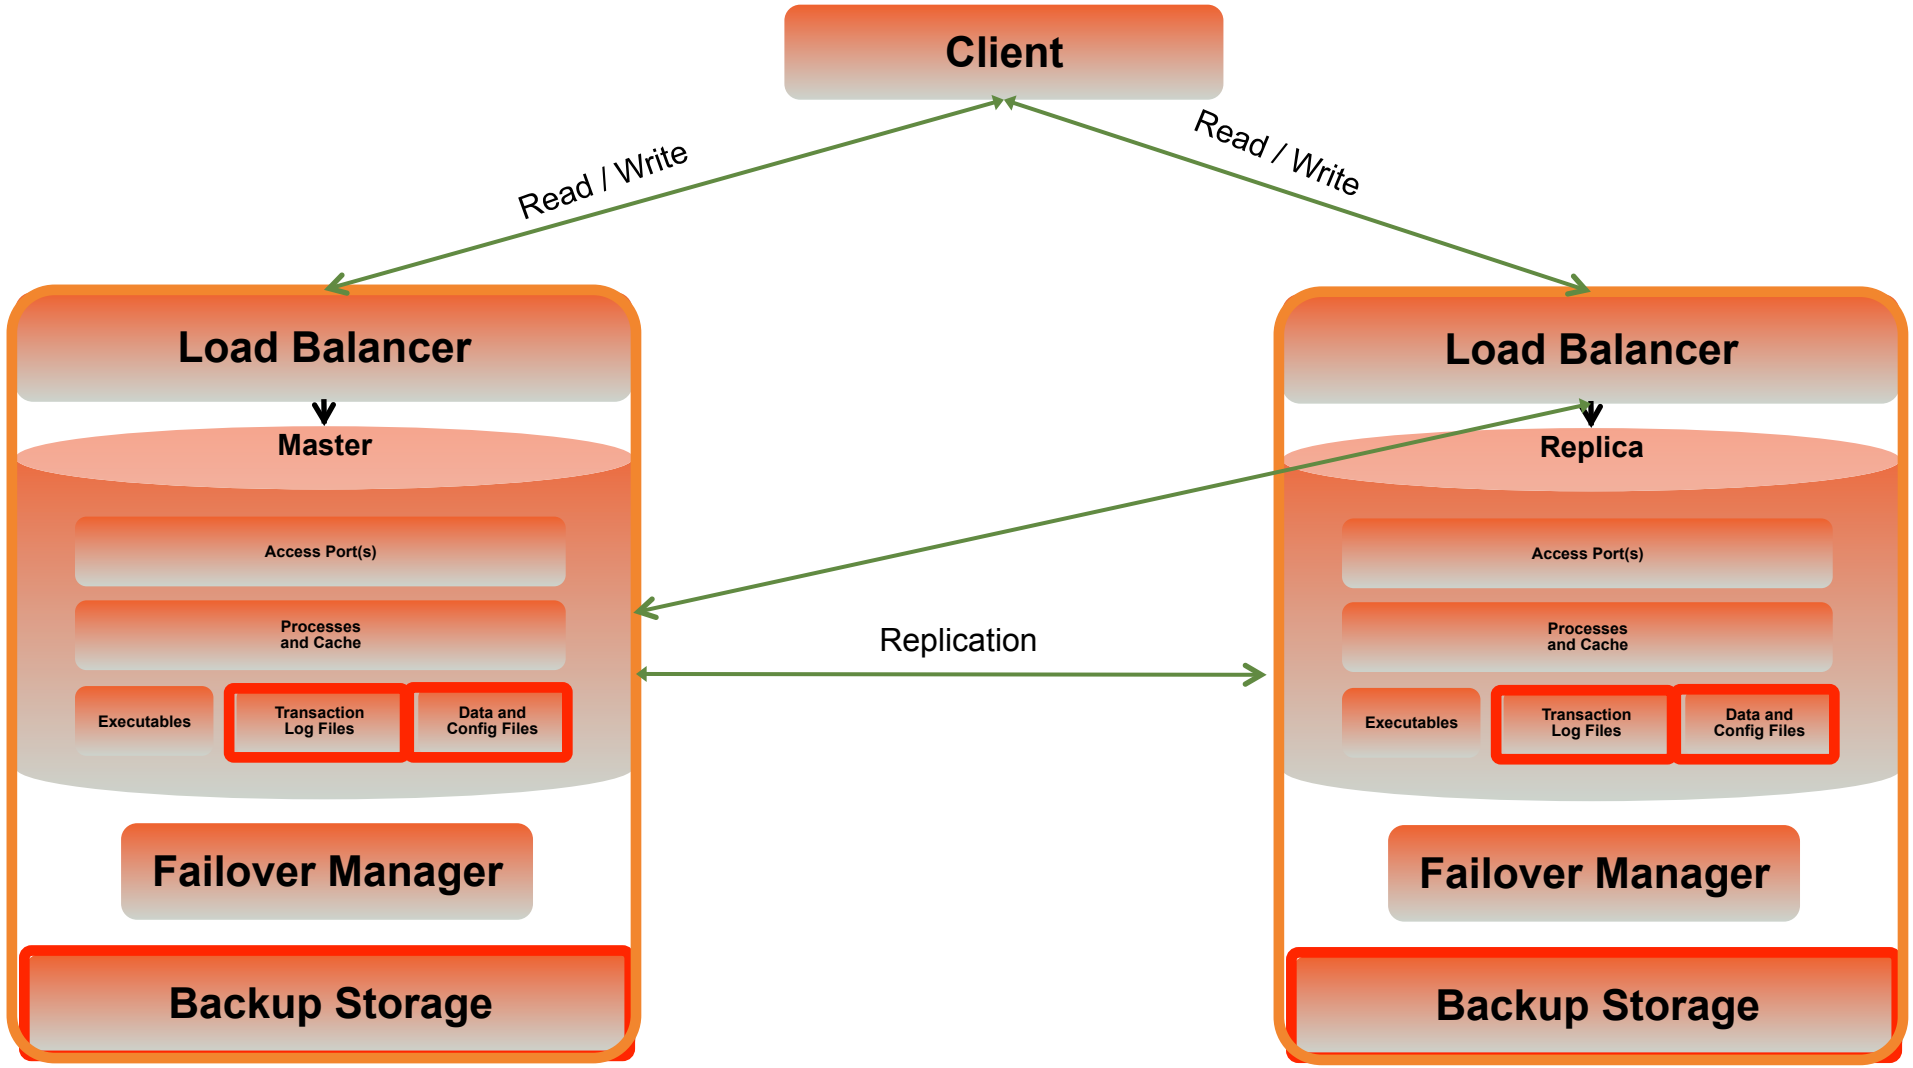

Automatic Failover

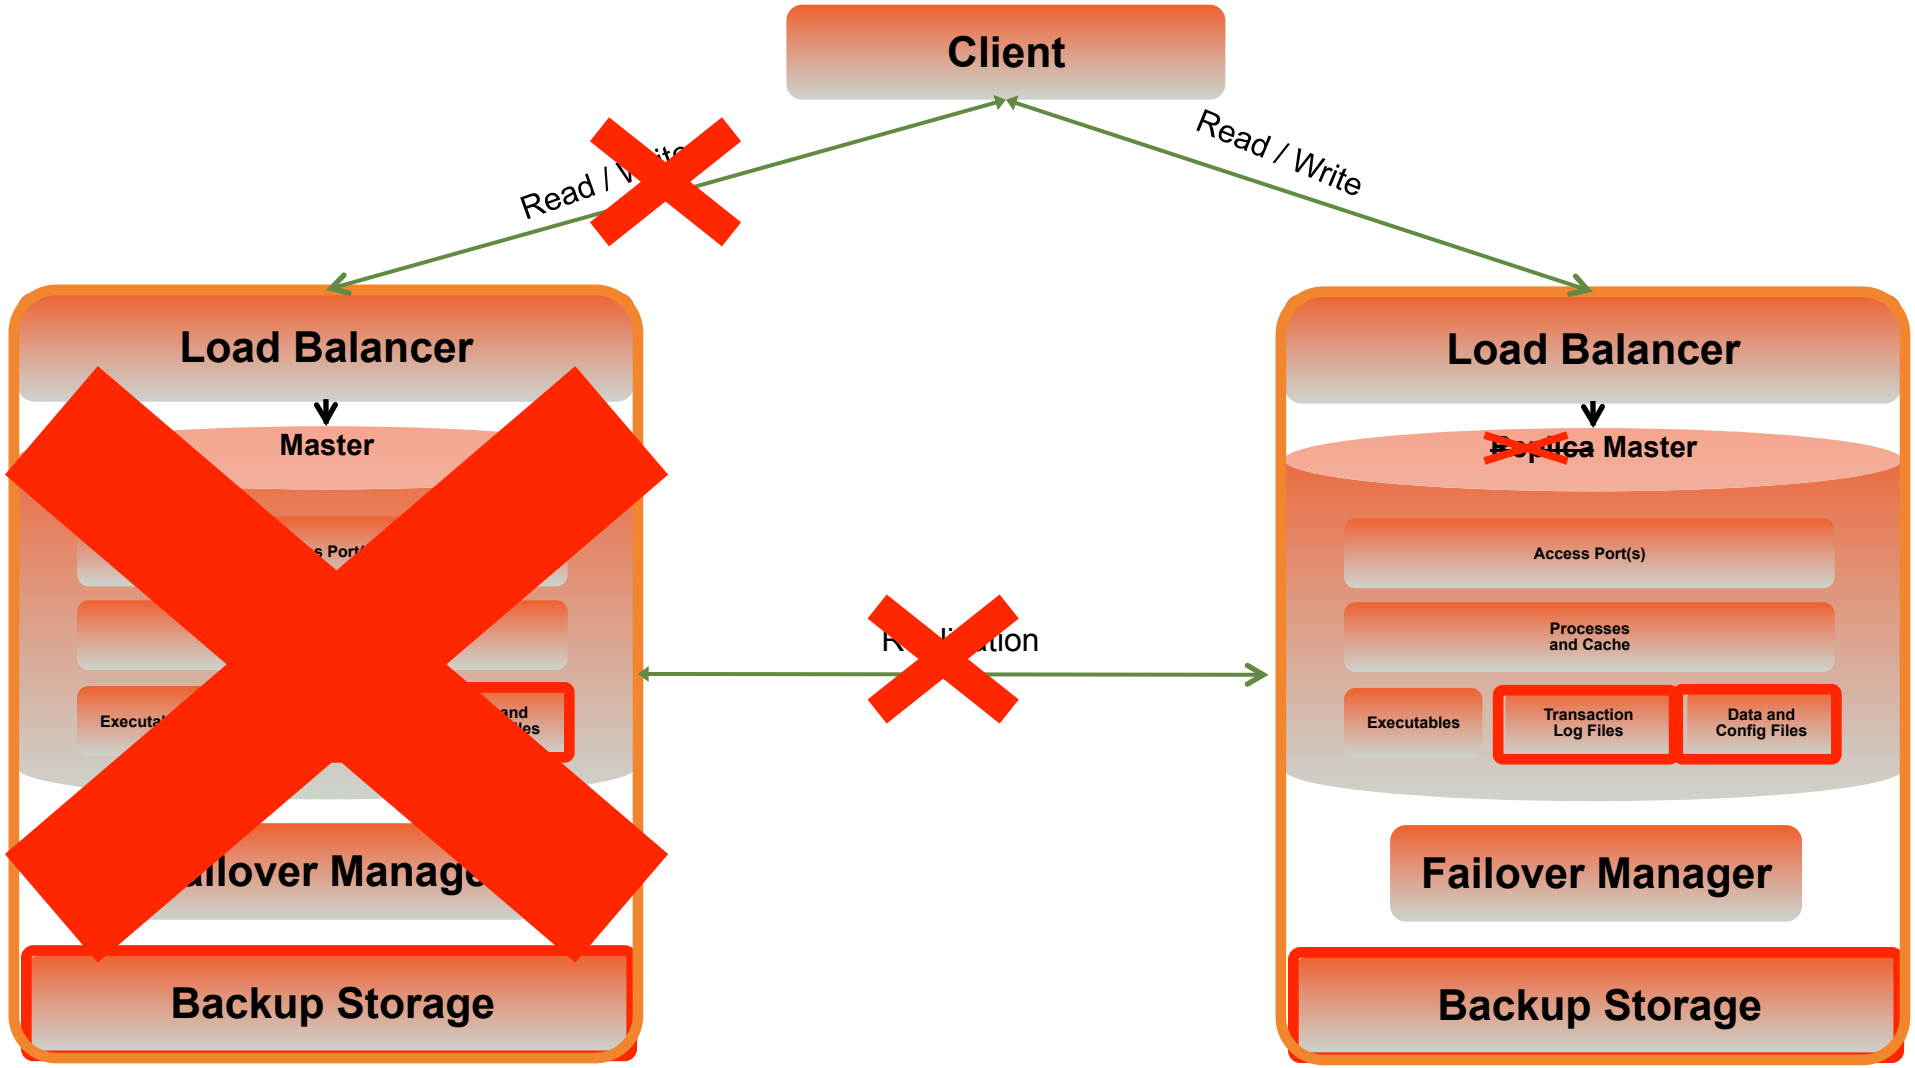

Smart Container Advantages

- Simple atomic containers, can be master or replica or both
  - Full HA environment has no more than 5 containers (4 DBs and 1 Backup)
- Clients access each container as if they are all masters, so leverage any load balancer
  - Kubernetes/OpenShift “Services”
  - HA Proxy
  - <insert your favorite TCP load balancer here>
- Scaling up/down is adding/removing a container in any environment
  - Kubernetes, OpenShift, Compose, Custom, etc.
- Containers can start from the latest backup rather than streaming full backup from master

Microservices

- Microservices are a “*particular way of designing software applications as suites of independently deployable services*” *
- Characteristics *
  - Componentization via Services
  - Organized around Business Capabilities
  - Products not Projects
  - Smart endpoints and dumb pipes
  - Decentralized Governance
  - Decentralized Data Management
  - Infrastructure Automation
  - Design for failure
  - Evolutionary Design

* http://martinfowler.com/articles/microservices.html

Microservices - Databases

- Same characteristics for databases in a microservices architecture
  - Componentization via Services
    - Access DB via an generic OpenShift service
  - Organized around Business Capabilities
    - Deploy DB in each OpenShift project
  - Smart endpoints and dumb pipes
    - Make the DB container smart, not the load balancer or client
  - Decentralized
    - Deploy DB in each OpenShift project
  - Design for failure
    - Make failover and restart a ‘normal’ thing
  - Evolutionary design
    - Incremental progress, don’t try to boil the ocean
    - Expect constant improvement, not upfront perfection

EDB Postgres for RED HAT® OPENSHIFT Container Platform

EDB POSTGRES™

OpenShift Platform-as-a-Service

- Platform-as-a-Service to support agile development
- Flexible provisioning from Docker container images
- Standards and workflow automation leveraging Kubernetes increase productivity

EDB Postgres for OpenShift Architecture

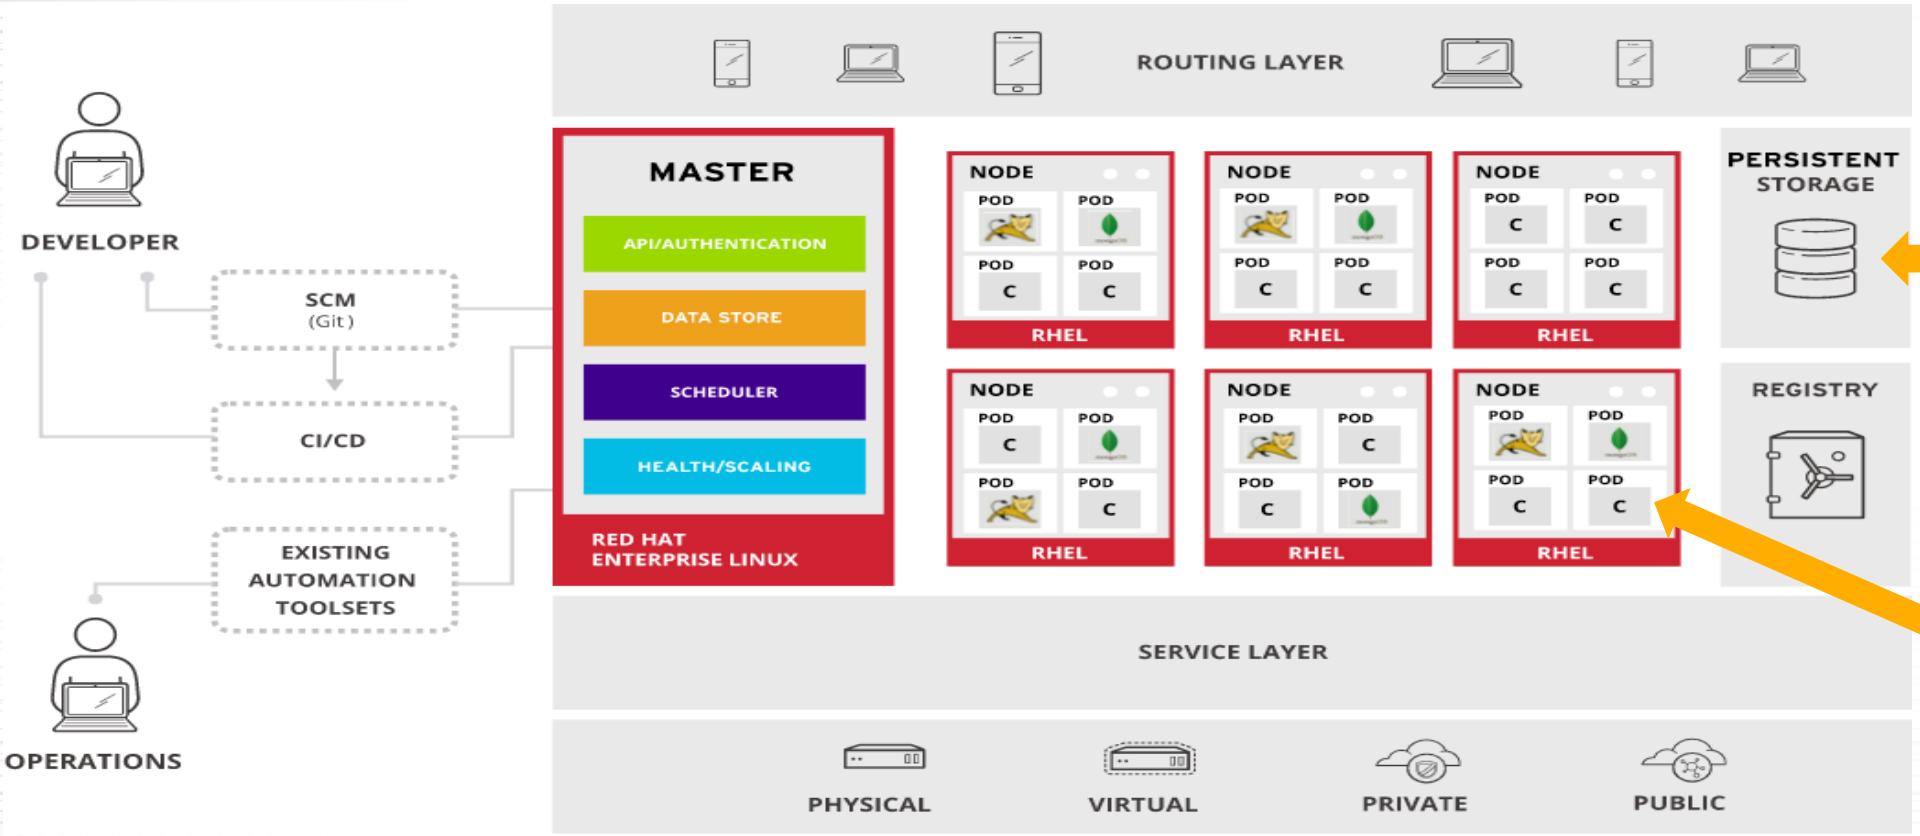

EDB Postgres configuration and data files

EDB Postgres deployed in container

Source: https://www.openshift.com/container-platform/features.html

EDB Postgres Integration for OpenShift

- Docker containers for the EDB Postgres Platform
- Kubernetes and YAML scripts for automation and orchestration
- DevOps self-service database provisioning through OpenShift's Management Console
- Developers can focus on the application rather than on the database
- DBAs retain visibility and control over the databases

EDB POSTGRES

Standardizing and automating EDB Postgres deployments through OpenShift

- EDB Postgres as persistent storage for OpenShift applications
- Agility in app development for your digital business
- Easy and fast provisioning with flexible scalability for optimal performance

EDB Postgres Containers

- Easy deployment of the EDB Postgres Platform

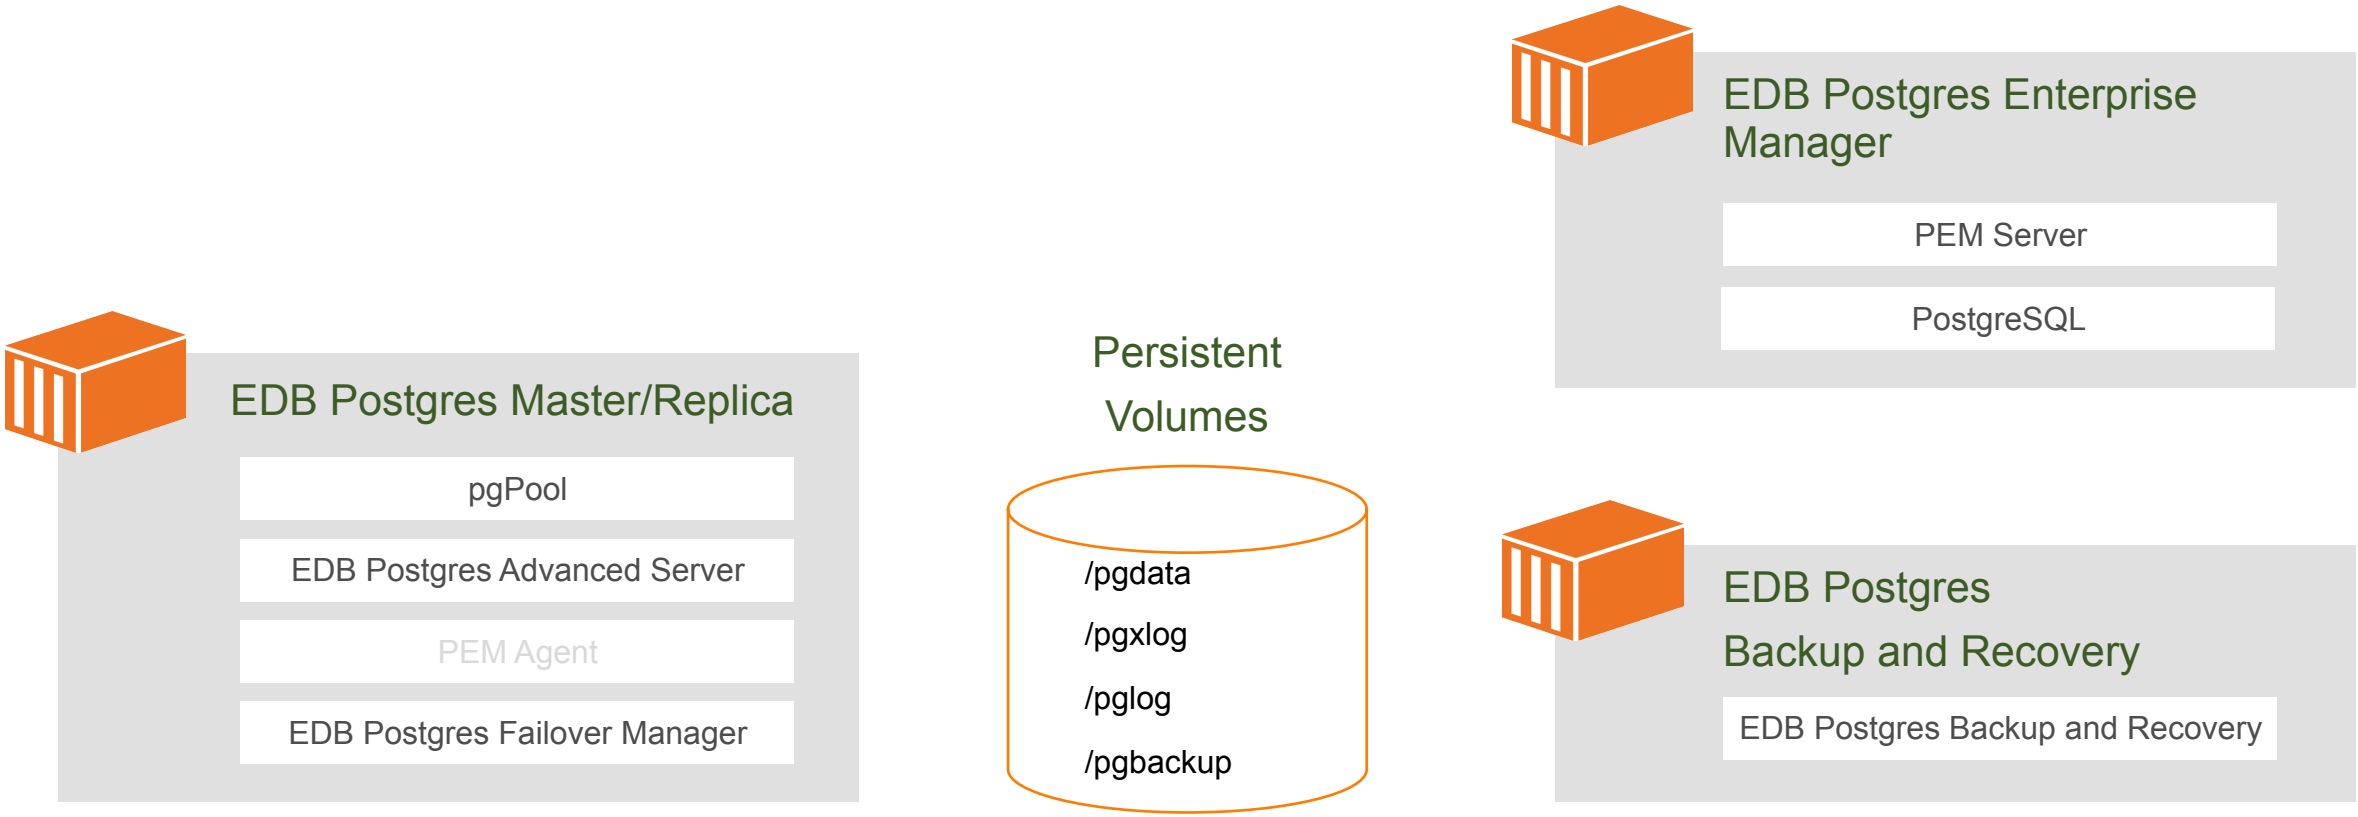

EDB Postgres:
Go fast & get where YOU WANT TO GO.

Approach digital transformation holistically.

Reduce complexity for greater long-term success.

Lead the revolution.

EDB™ POSTGRES


EDB
POSTGRES
EDB POSTGRES CONTAINERS AND INTEGRATION WITH OPENSHIFT:
HTTPS://GET.ENTERPRISEDB.COM/DOCS/EDB_POSTGRES_CONTAINERS_V2.1.PDF?_GA=2.16899897.1477532771.1528121425-326313351.1523434668

Questions?


EDB
POSTGRES
THANK YOU
Ossi Karjalainen


EDB
POSTGRES
EDB AS A COMPANY

EDB POSTGRES™

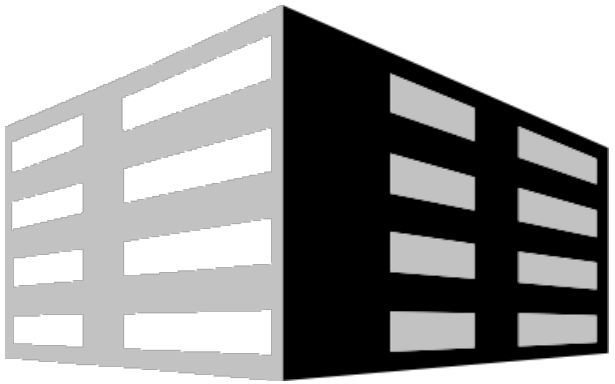

Founded in 2004

- **The world leader in open-source based Postgres software and services.**
- **Customer base > 4000**
- **300+ employees**
- **Offices worldwide**
- **Recognized RDBMS leader**
  - **Gartner and Forrester**
- **Recognized PostgreSQL community leader**

EDB POSTGRES™

Gartner Magic Quadrant and EnterpriseDB

EDB Recognized 5 Years In A Row on Gartner's Magic Quadrant

Figure 1. Magic Quadrant for Operational Database Management Systems

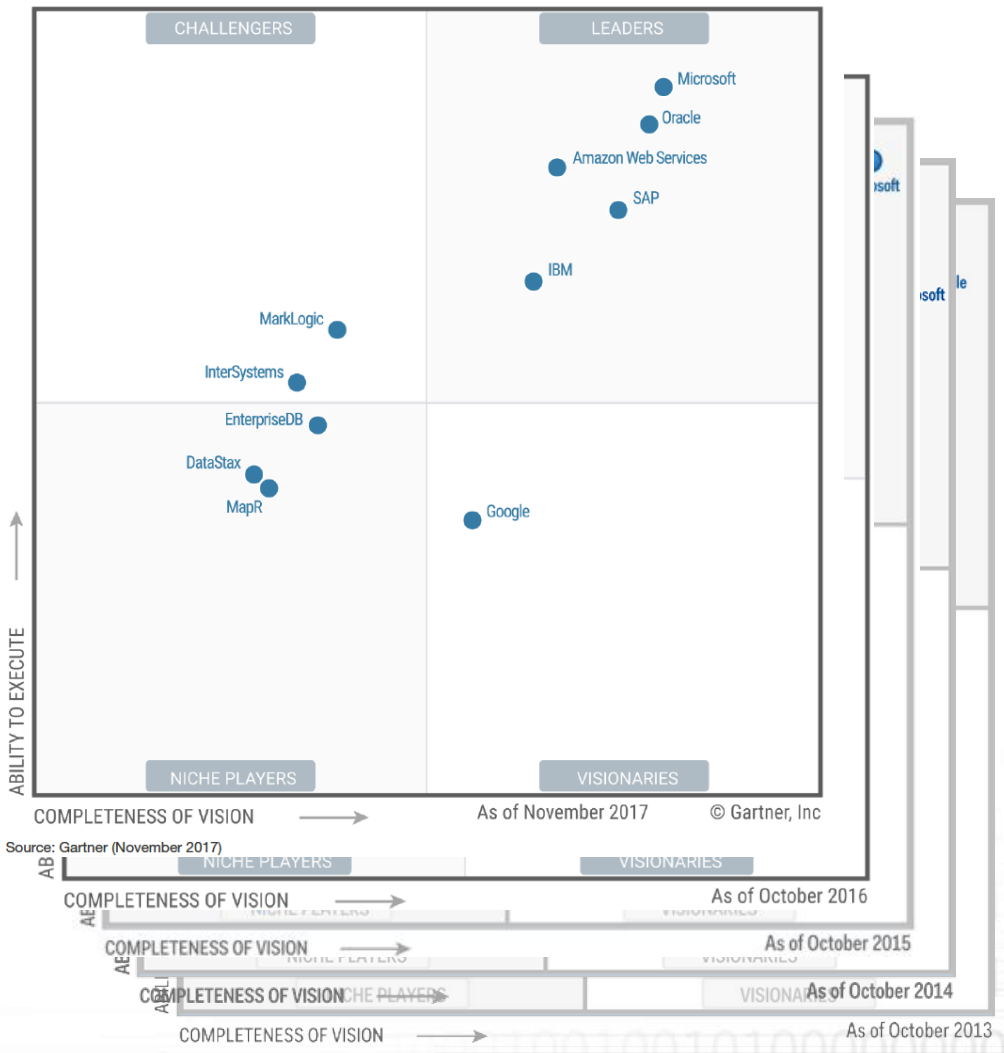

This graphic was published by Gartner, Inc. as part of a larger research document and should be evaluated in the context of the entire document. The Gartner document is available upon request from EnterpriseDB.

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research Publications consist of the opinions of Gartner's research Organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of Merchantability or fitness for a particular purpose.

The PostgreSQL Community

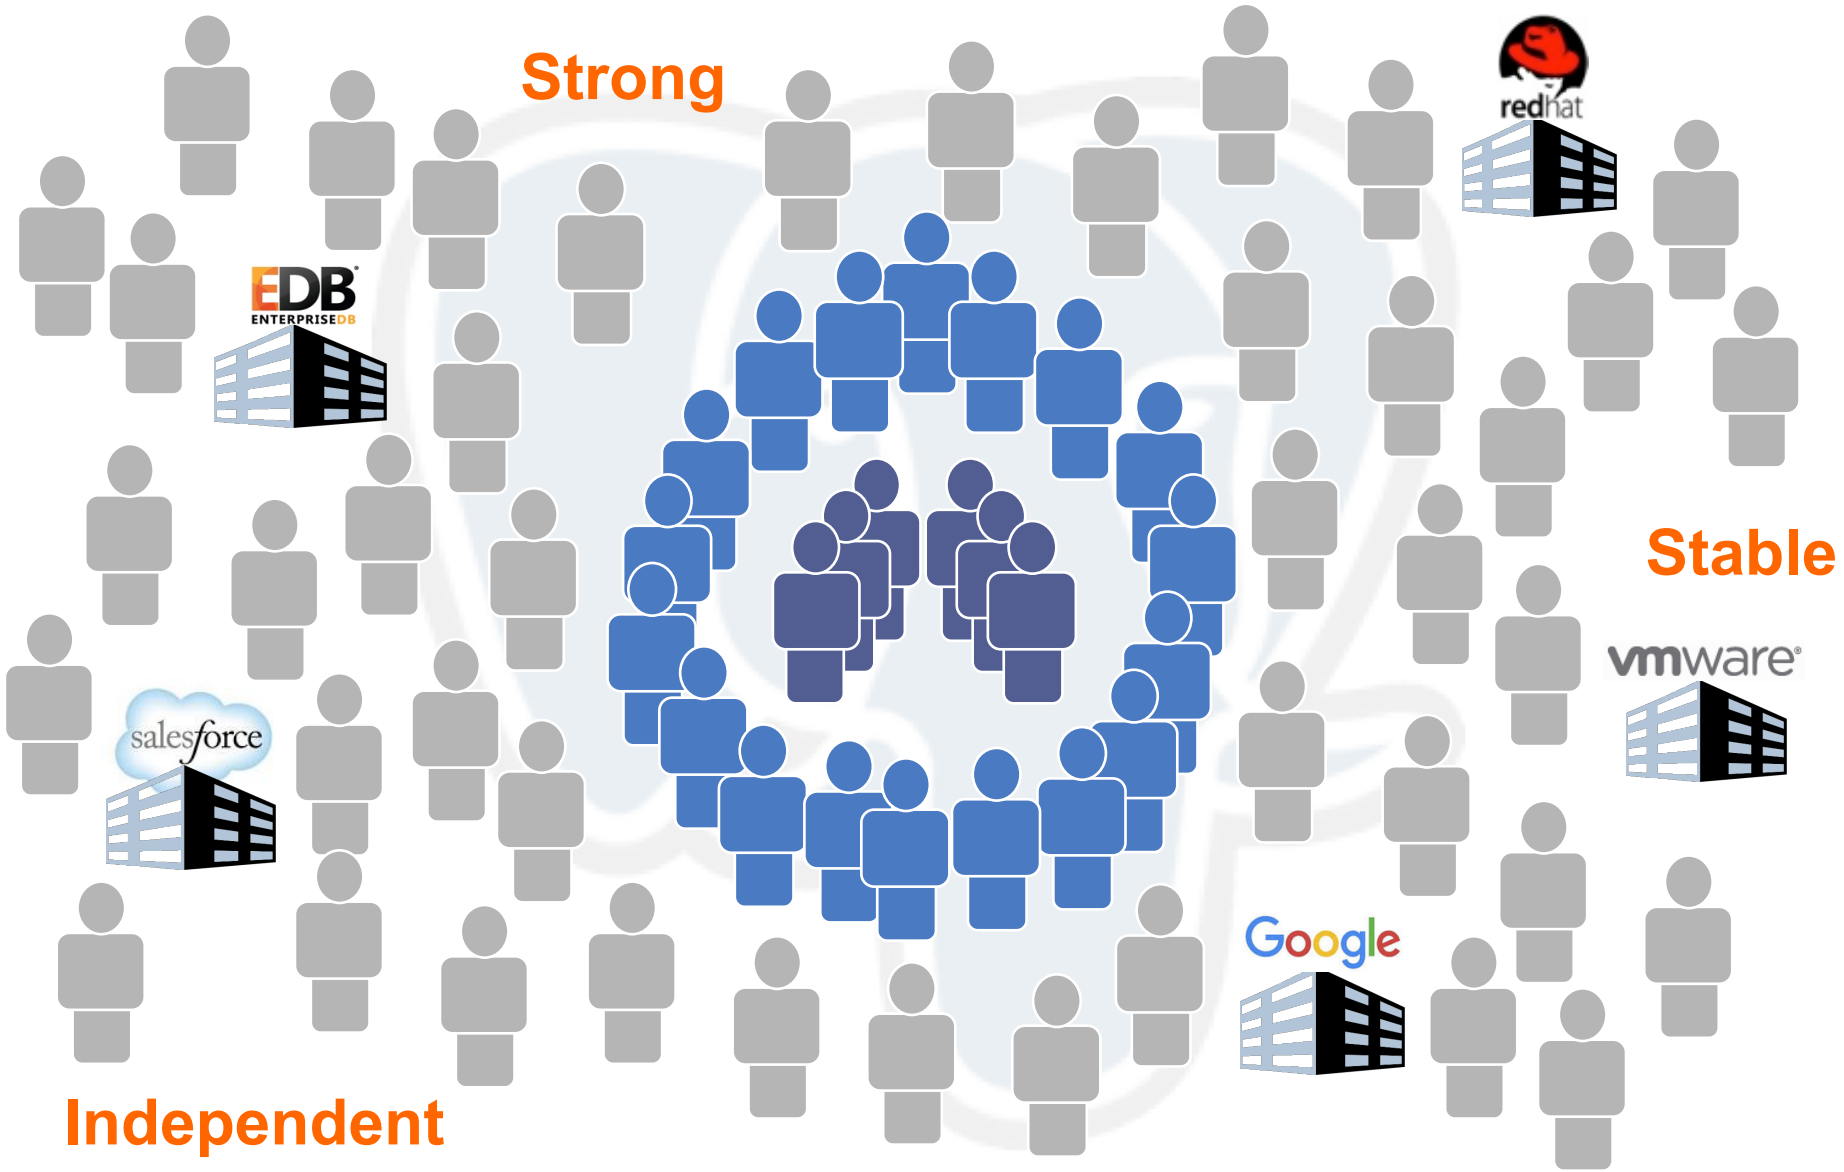

EDB POSTGRES

POSTGRESQL
innovation

ENTERPRISE
reliability

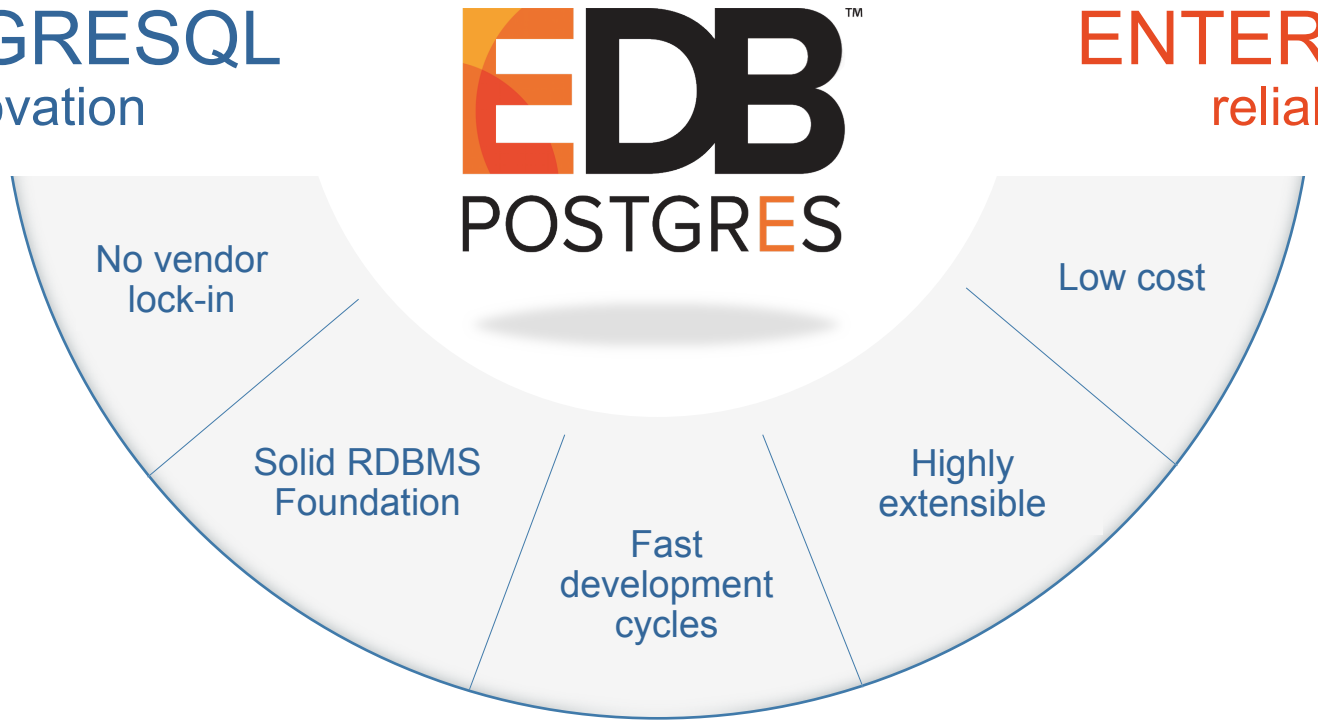

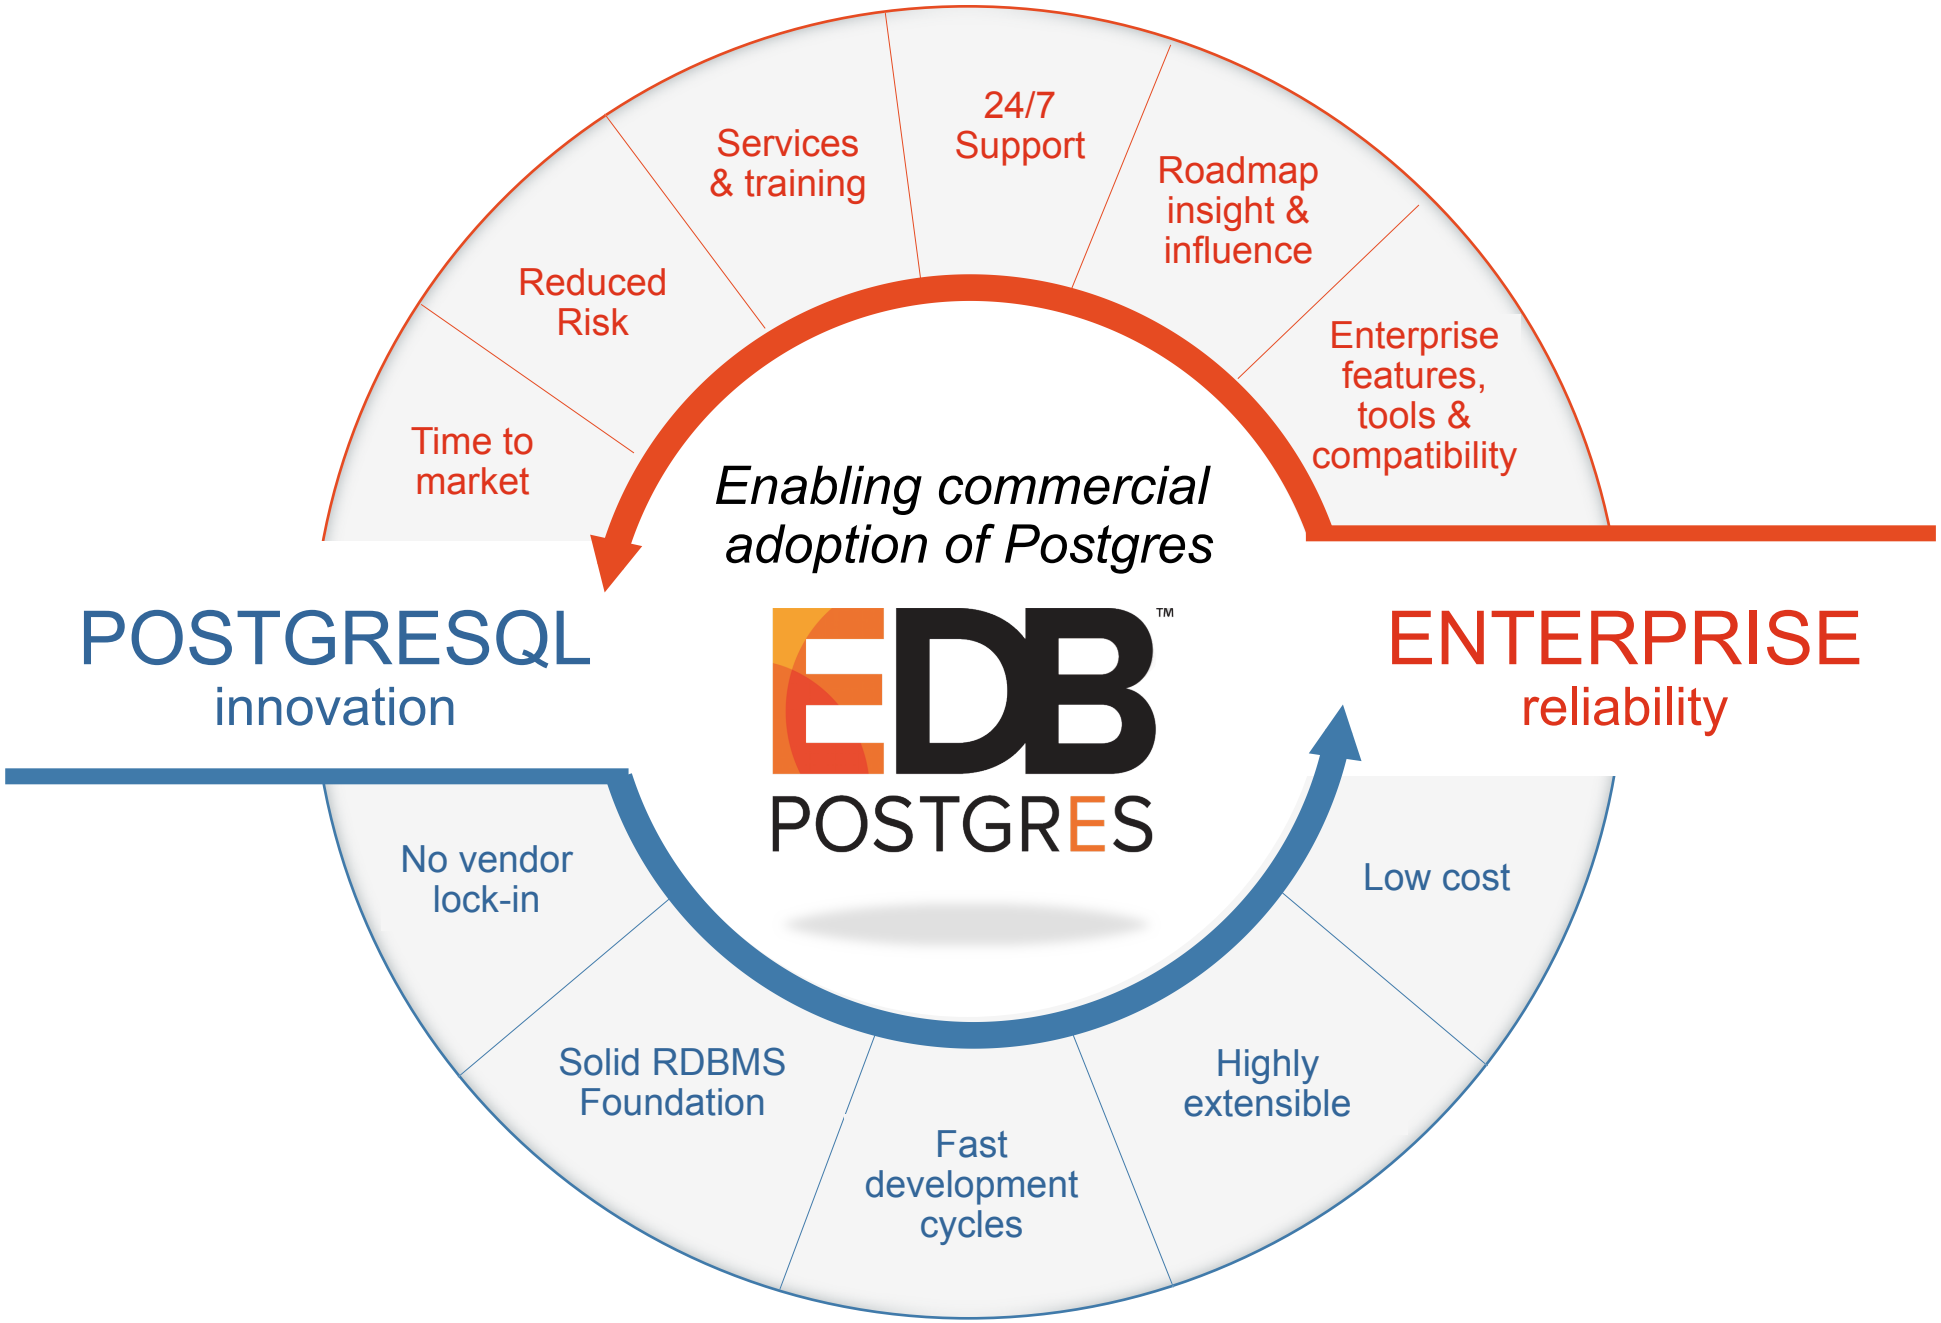
Time to market
Reduced Risk
Services & training
24/7 Support
Roadmap insight & influence
Enterprise features, tools & compatibility
Enabling commercial adoption of Postgres
POSTGRESQL
innovation
ENTERPRISE
reliability
EDB
POSTGRES
No vendor lock-in
Solid RDBMS Foundation
Fast development cycles
Highly extensible
Low cost

Postgres

Advanced Server

EDB POSTGRES

Database Compatibility for Oracle

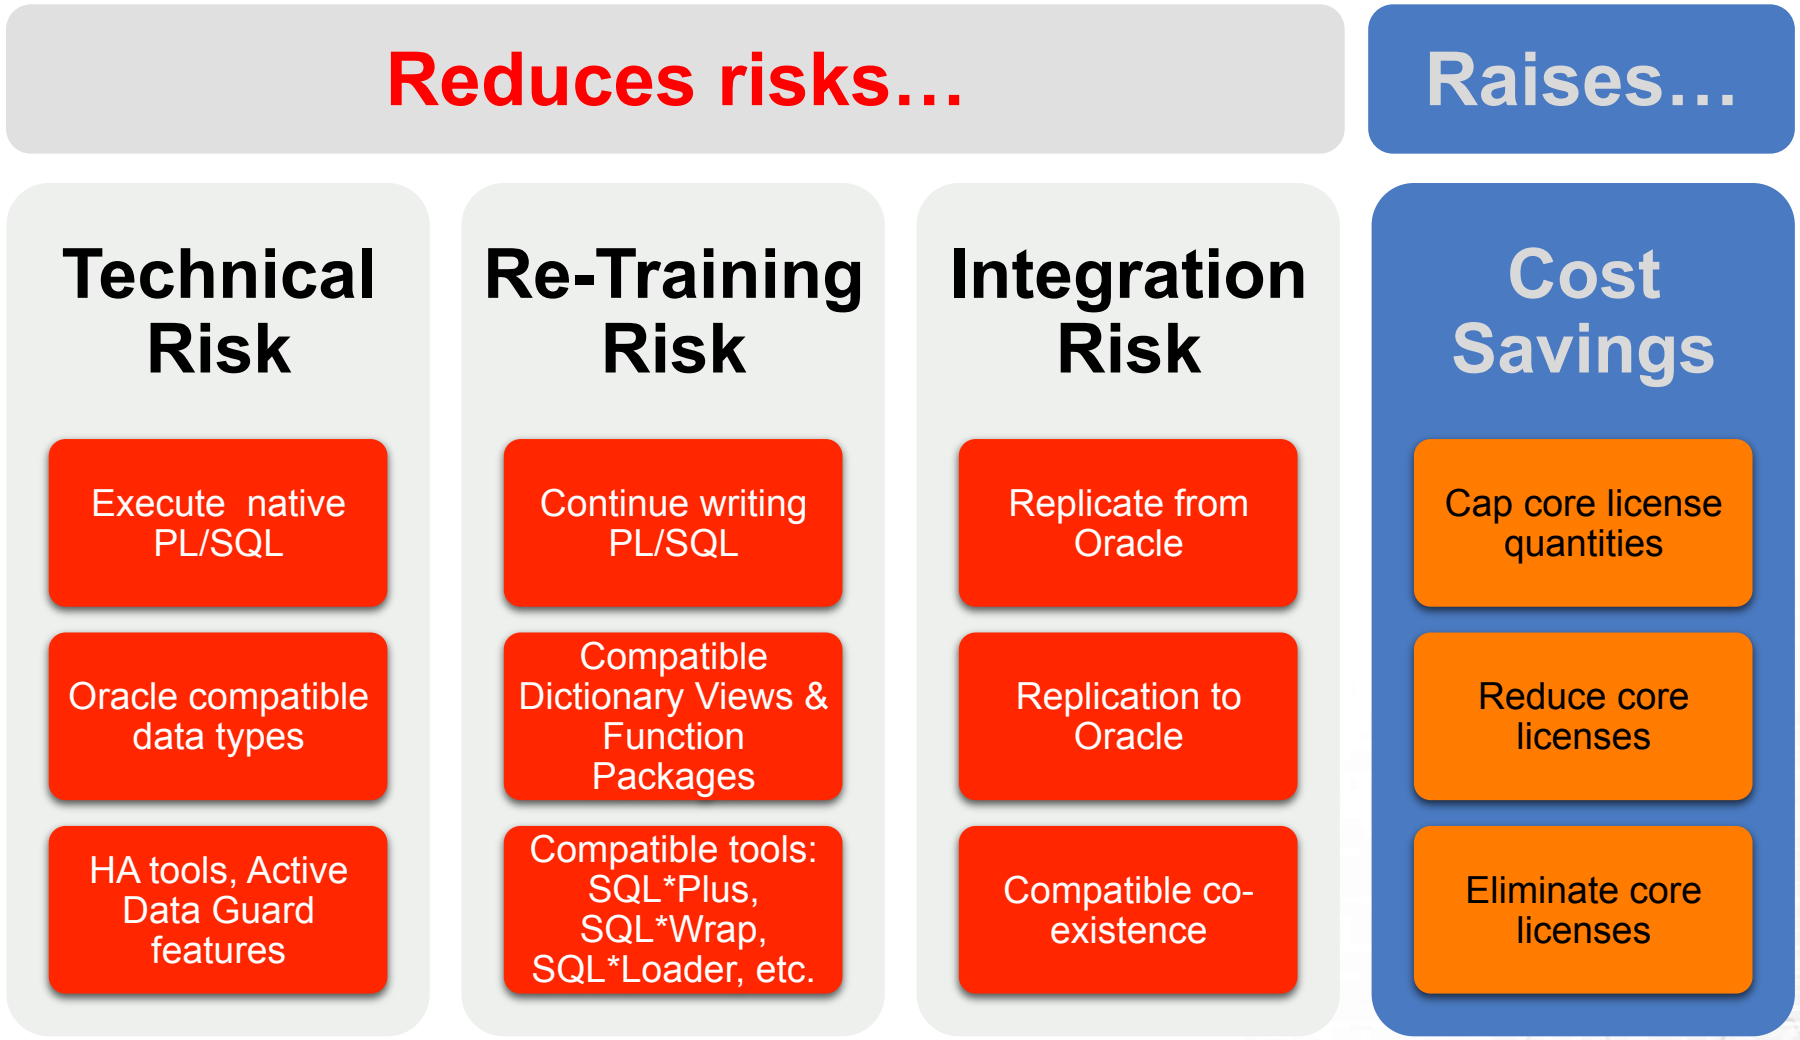

EDB Postgres Platform 2018

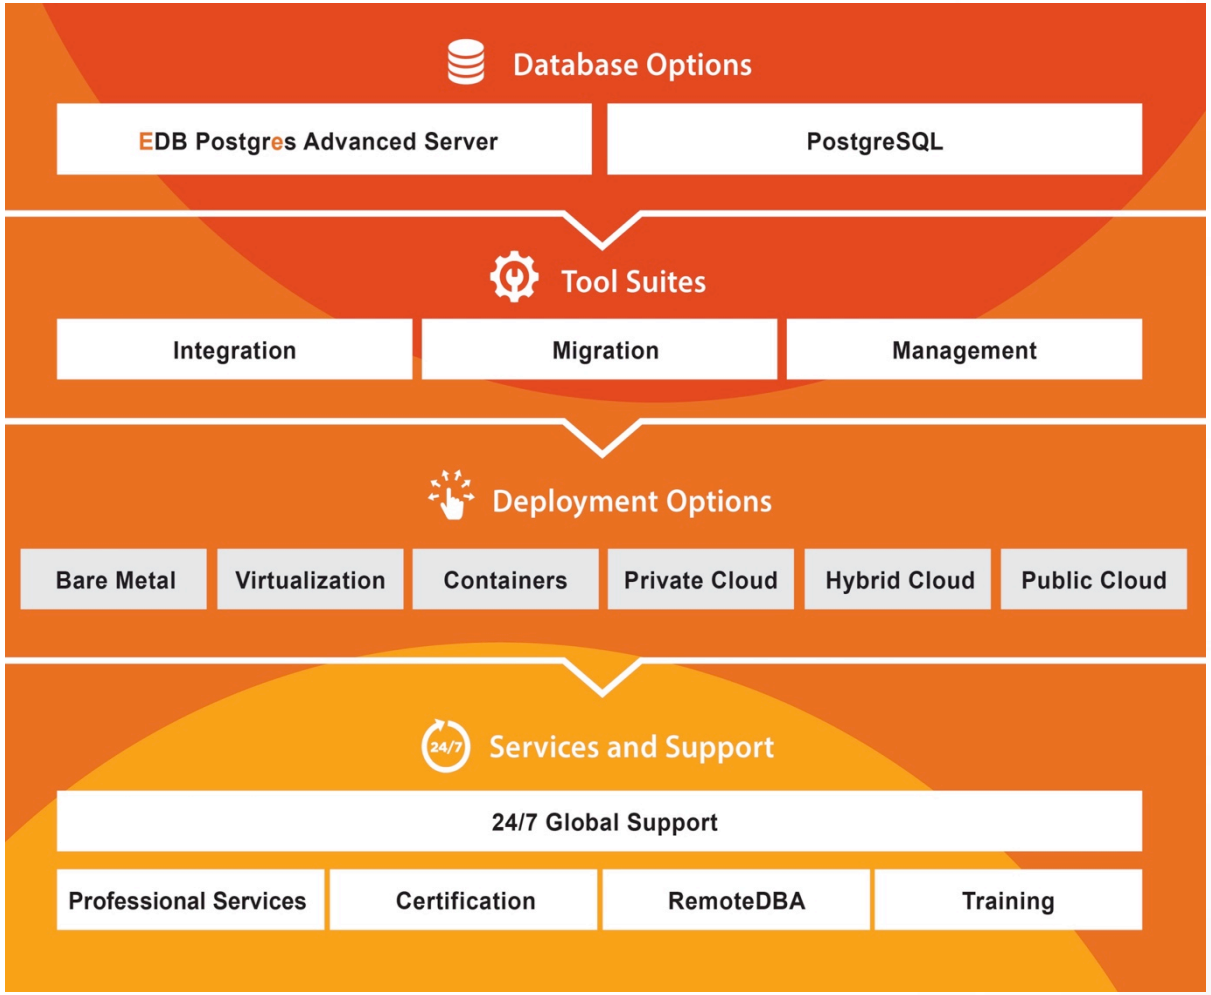

EDB POSTGRES

Cloud Management Flexibility

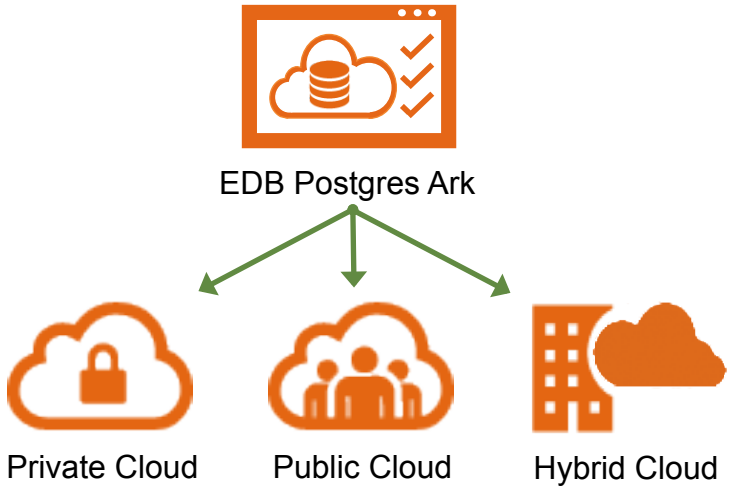

Low TCO & Big ROI

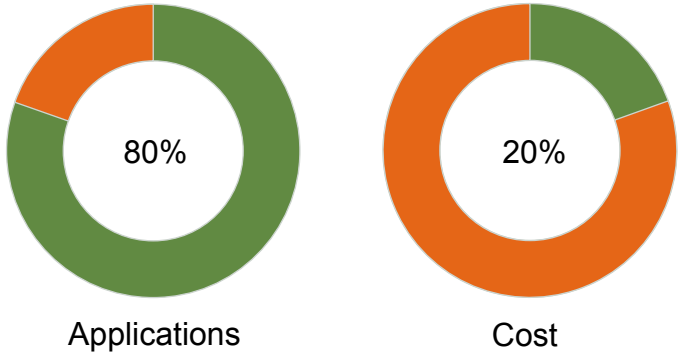

Enterprise Postgres

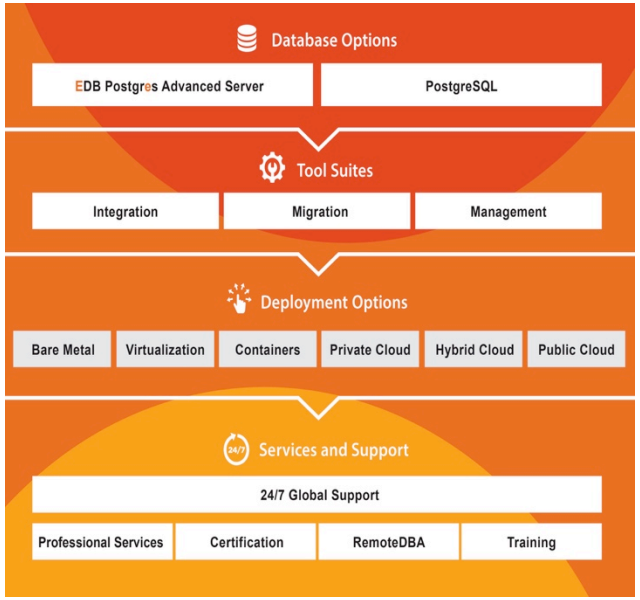

DBMS for Digital Business

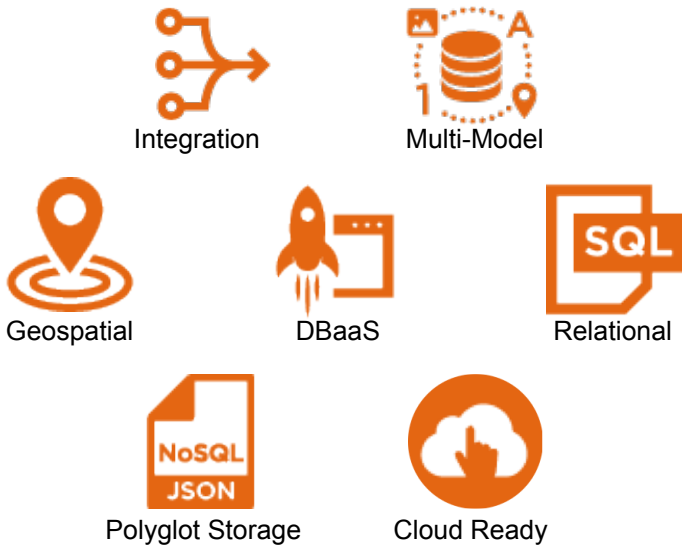

EDB Capabilities

- Innovation
  - Advanced Server: Great database
  - Faster and more scalable every year
  - DevOps support
- Reliability
  - EFM, BART, xDB, …
- Legacy Leverage
  - Oracle compatibility
  - Migration tool kit

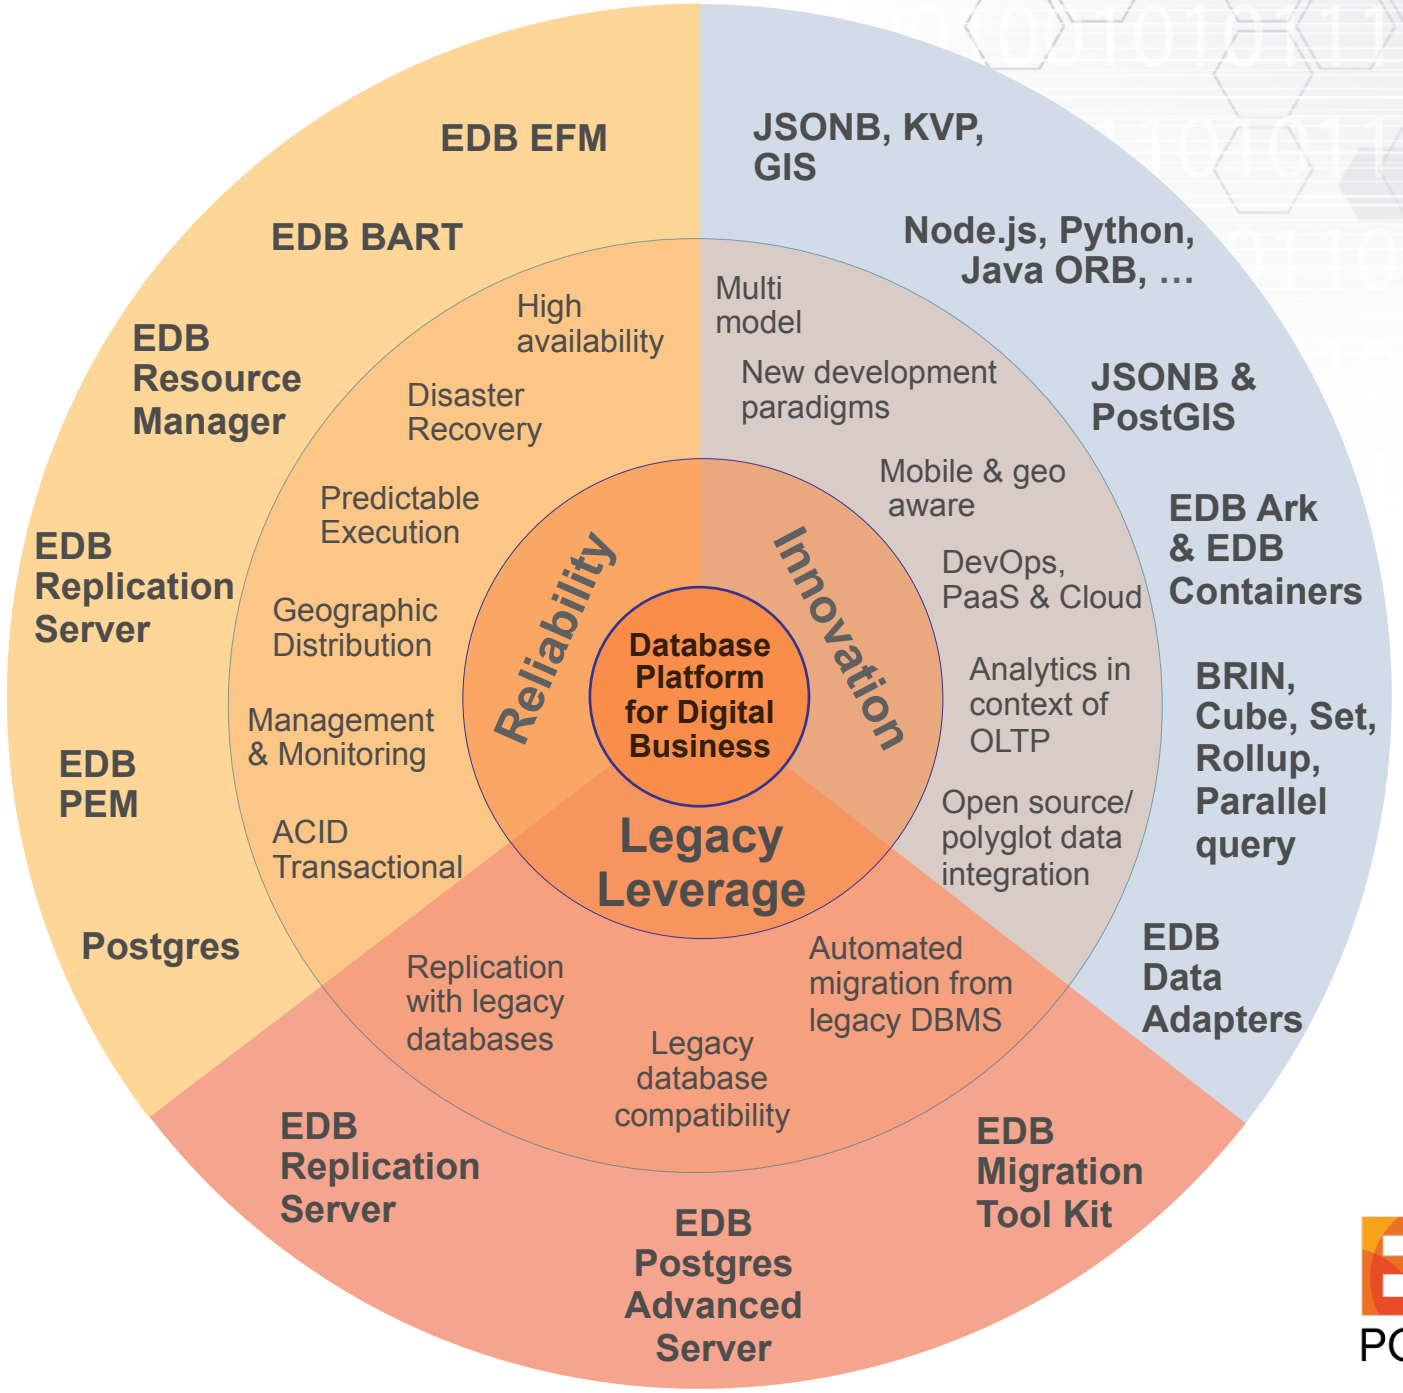

EDB Customers

EDB currently has over 4,000 total customers including 85 of the Fortune 500 and 169 of the Forbes Global 2000

EDB POSTGRES