

Containing Without Containers

Robert Tracey

Cognitive and HPC Cloud Developer/System
Administrator

IBM Research

- 3,000+ (researchers and engineers)
- 5 Nobel Prizes
- 6 Turing Awards
- 10 National Medals of Technology
- 5 National Medals of Science
- Averaging 9.3 patents per day
- 22 years of patent leadership
- Collaboration with clients and universities
- Multi-disciplinary projects that lead to prototypes or long-term projects

12 labs. 6 continents.



Africa
Nairobi, Kenya

Australia
Melbourne

Haifa
Haifa, Israel

Tokyo
Tokyo and Shin-kawasaki, Japan

Almaden
San Jose, California, U.S.A.

Brazil
São Paulo and Rio de Janeiro

India
Delhi and Bengaluru

Watson
New York and Massachusetts, U.S.A.

Austin
Austin, Texas, U.S.A.

China
Beijing and Shanghai

Ireland
Dublin

Zürich
Rüschlikon, Switzerland

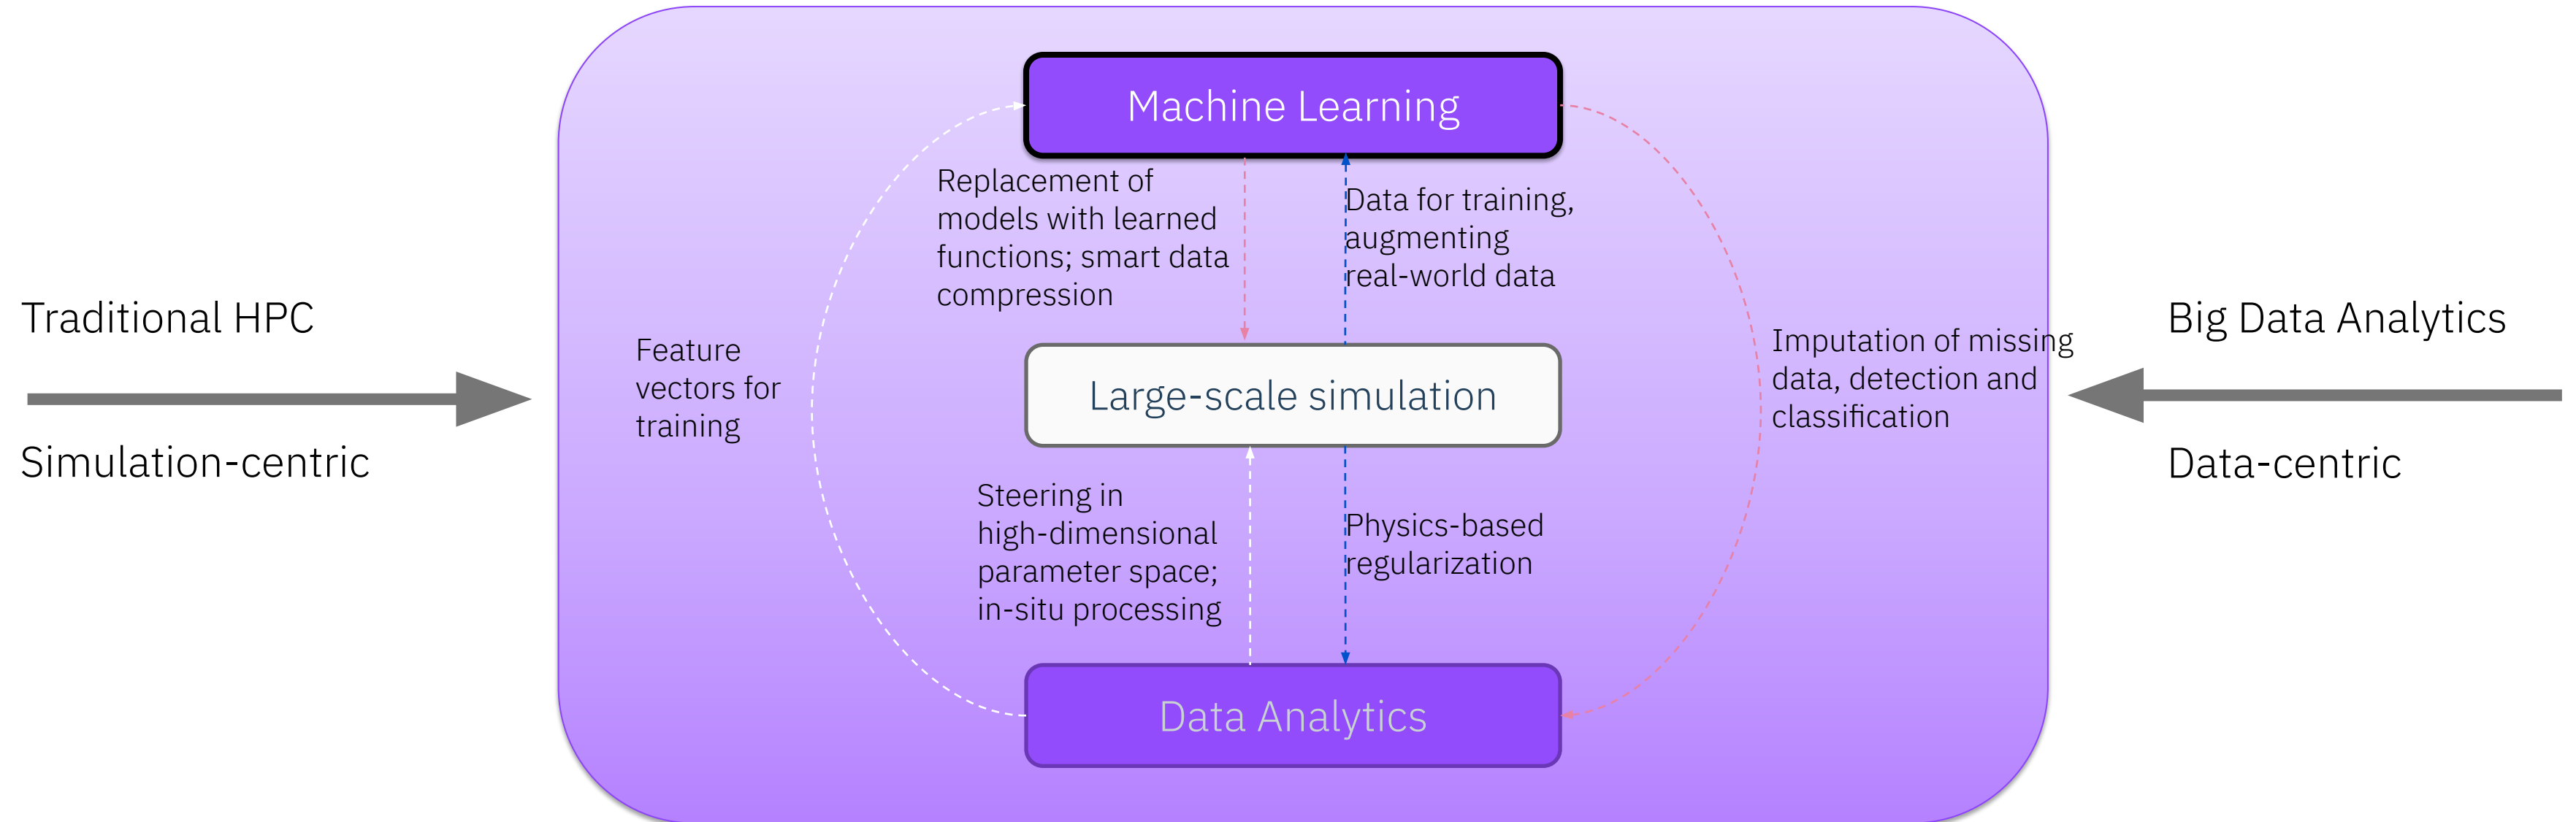
Designing Next-Generation Systems

Data-centric approach



- DoE's CORAL systems, ORNL's Summit and LLNL's Sierra, first instantiation of IBM's data-centric approach for system design
- Platforms for innovative emerging workflows, laying the groundwork for efficient scientific discovery at Exascale
- Designed for data and AI from the ground up
- Summit, currently world's most powerful system
- 5 2018 Gordon Bell finalists used Summit or Sierra

Era of Data-Centric and Intelligent Discovery



- Explosion of data generated by large-scale simulation leading to a paradigm shift: from simulation-centric to **data-centric discovery**
- Data analytics and machine learning used to turn reams of simulation data into **actionable** information that can be used for better interpretation and steering
- Applying machine learning for making existing simulation codes more **intelligent**, more productive, and more robust
- Increasing interest in **large-scale** analytics and machine learning on high-end platforms
- Emerging hybrid workflows that embody the **entire** inference cycle of discovery
- Co-deployment of **heterogenous** software stacks

What's wrong with Containers?

Simple answer: Nothing!

They provide:

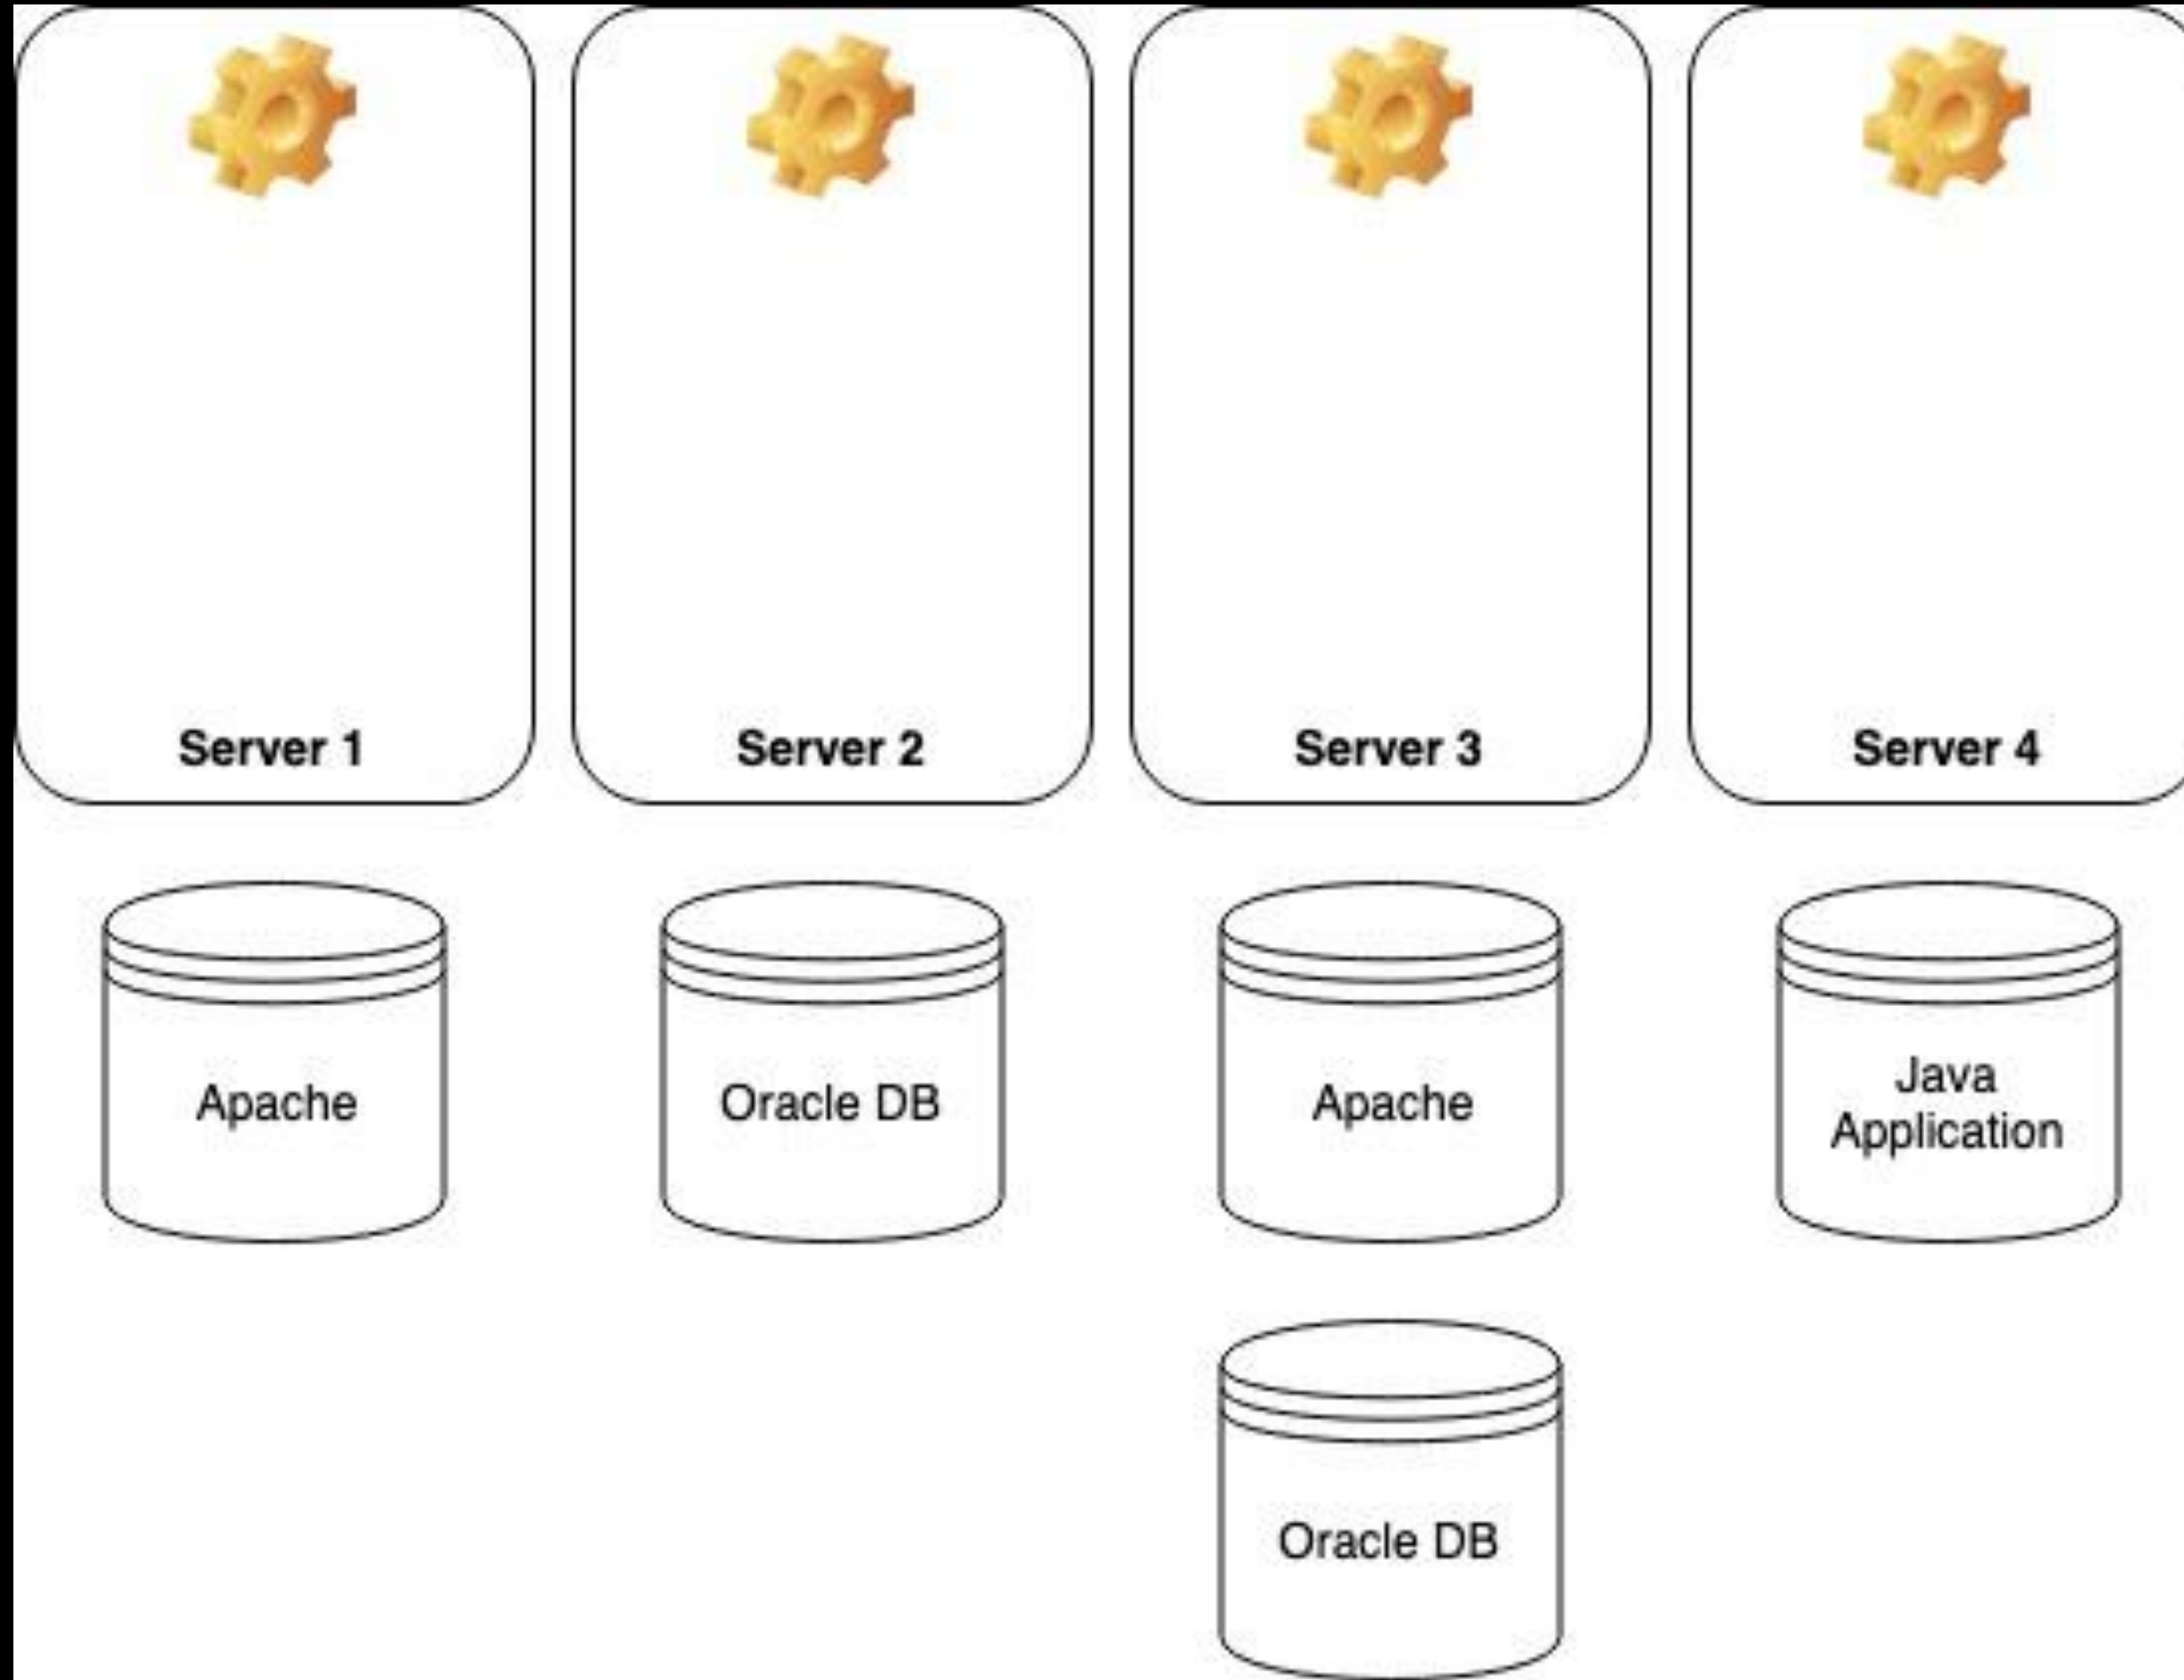
- Scaling
- Isolation
- Platform Independence
- Simplify Pipelines for developers
- And many more reasons

Then why are you here?

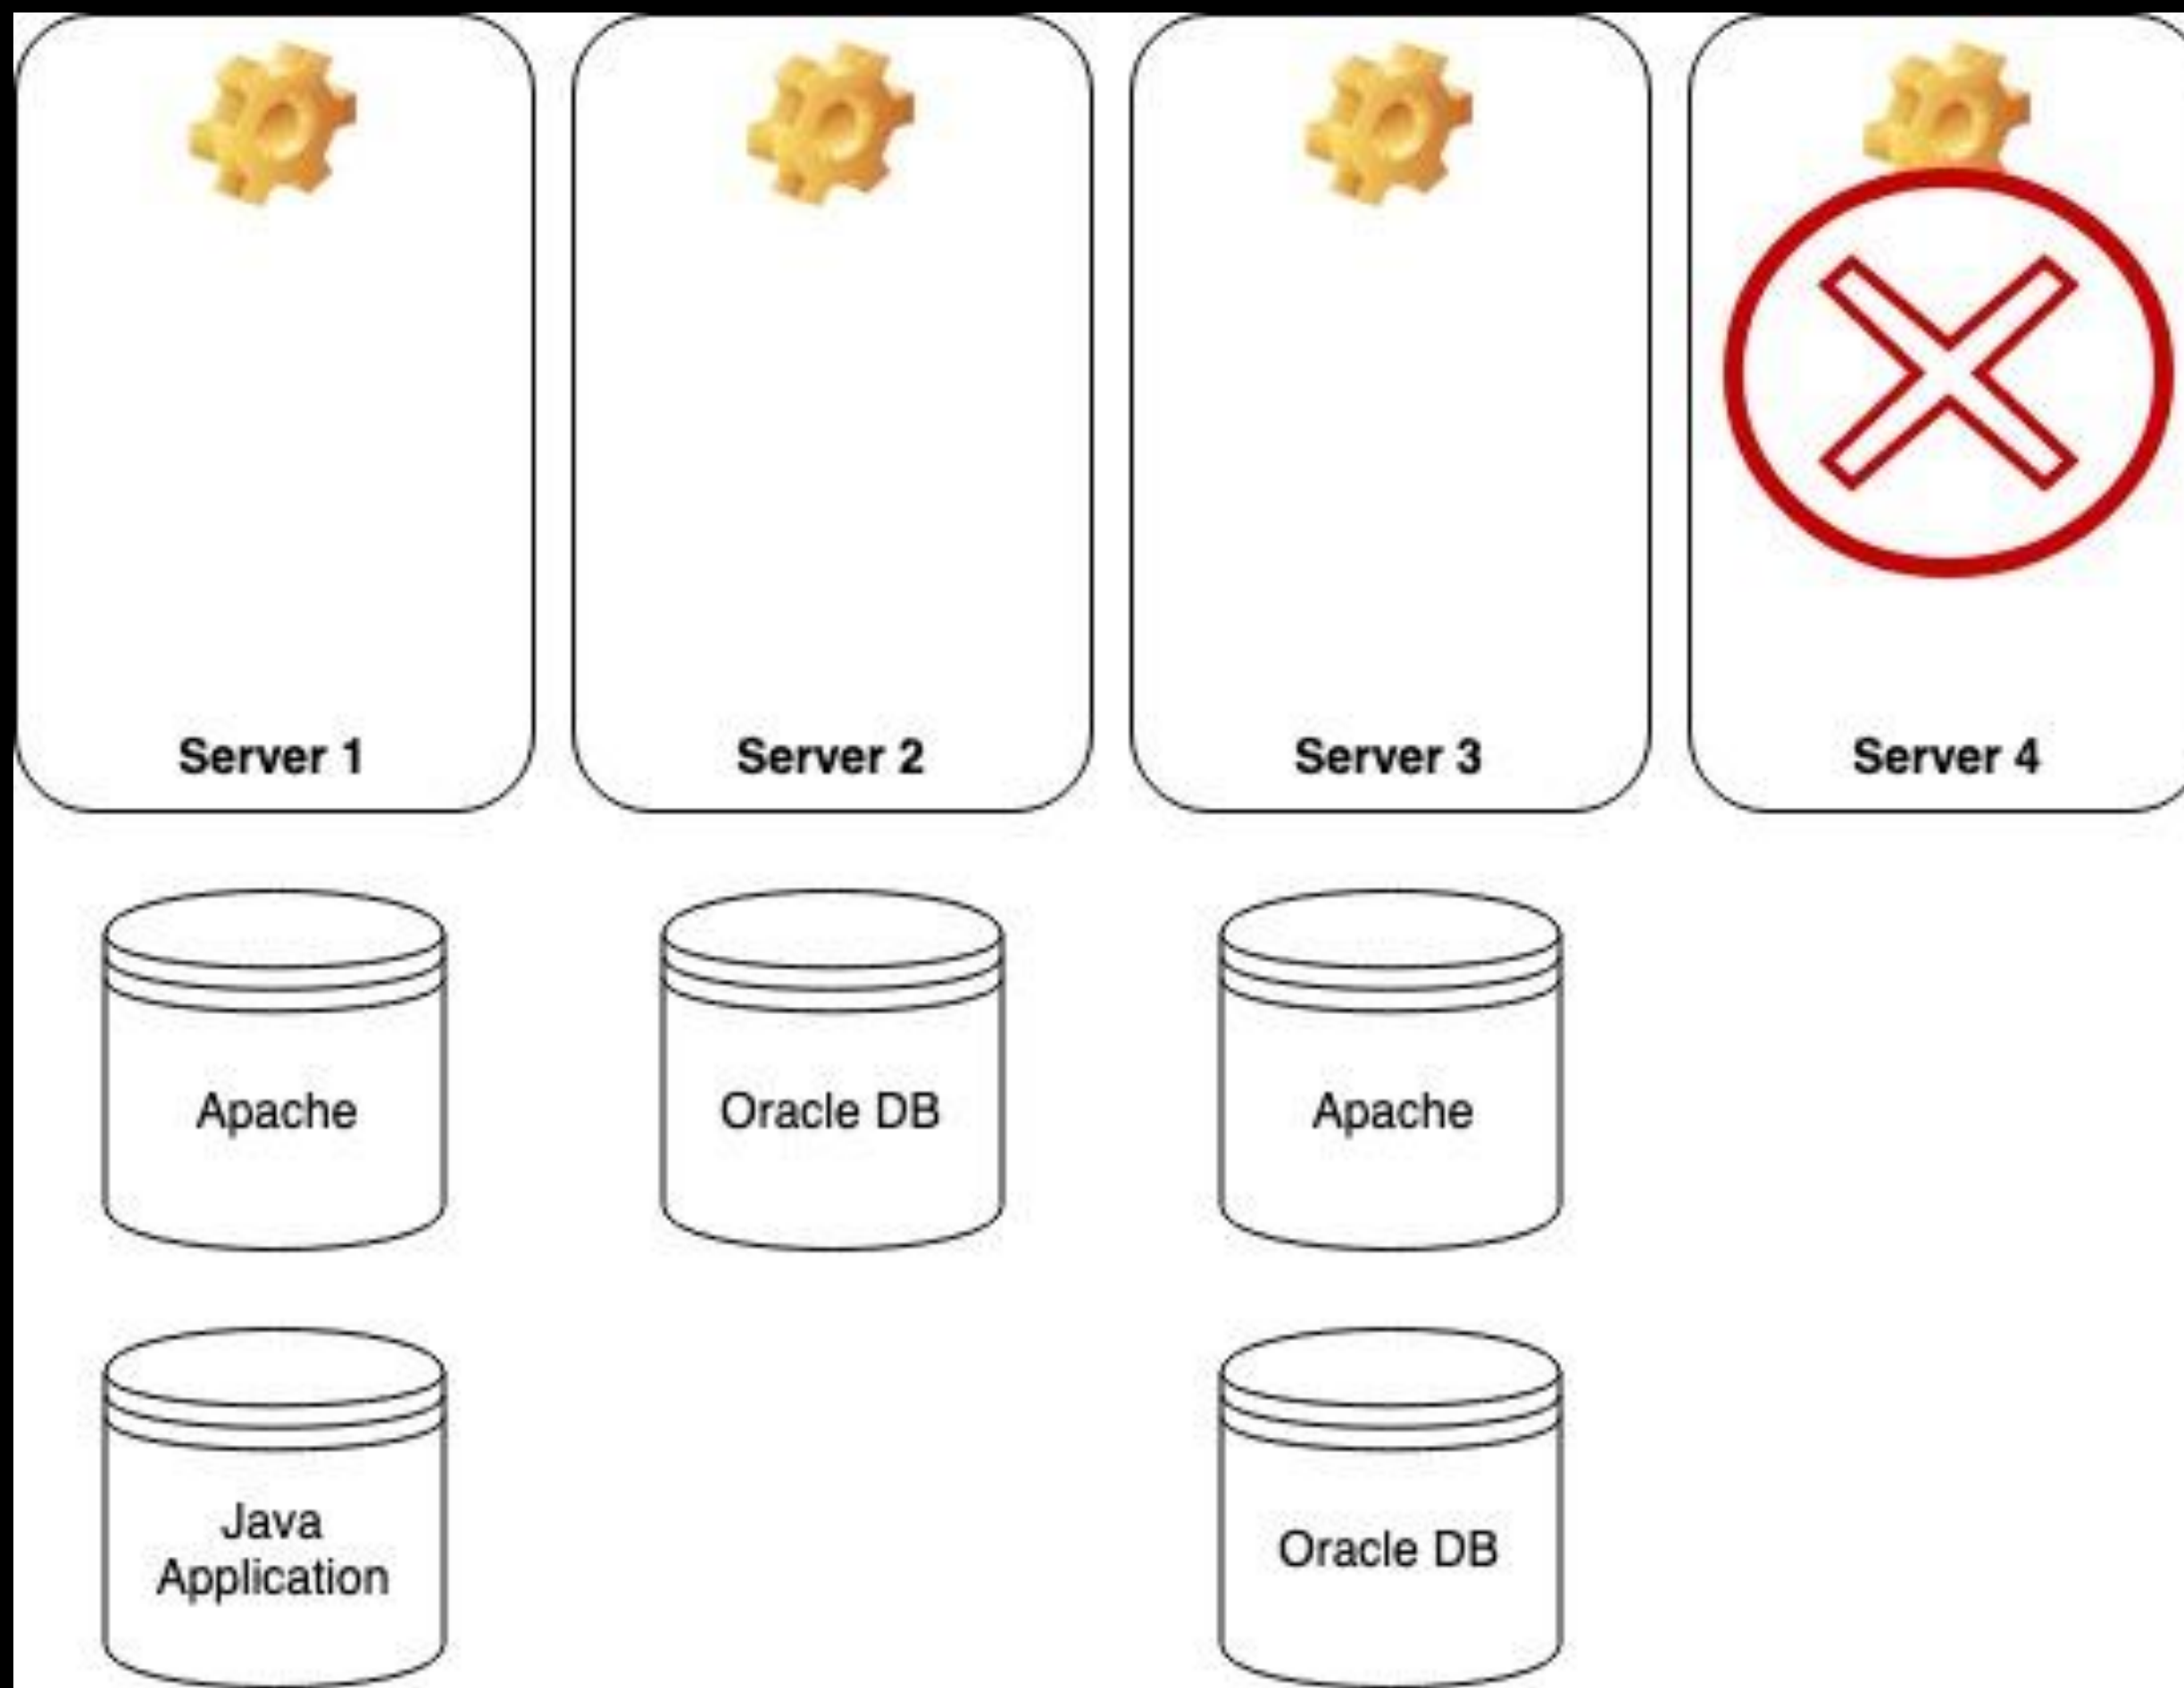
As good as containers are they have some downsides:

- Security
- Fractured eco-system
- Overhead
- Extra software to install

Lets look at HA Environments



Lets look at HA Environments



So how is this done in HA?

Two main ways:

- To install the applications and libraries on every server
- To contain all the needed code, libraries and binaries on a each disk

So does this work and for what?

I used two pieces of software to test this, both work very differently

1. MongoDB

- Useful for ML and Big Data workloads as schema-less
- Tunable
- Prebuilt binaries

2. Apache

- Useful for viewing results
- Lots of addons and highly configurable
- Requires manually compiling

Lets have a look at some code

```
#!/bin/bash
#BSUB -J ls-test
#BSUB -oo stdout.txt
#BSUB -eo stderr.txt
#BSUB -q panthel
#BSUB -W 00:25
#BSUB -n 1
```

```
LOGDIR=/gpfs/panther/scratch/rob/mongodb-linux-ppc64le-enterprise-rhel71-4.1.4/log
DBDIR=/gpfs/panther/scratch/rob/mongodb-linux-ppc64le-enterprise-rhel71-4.1.4/data/db
EXECDIR=/gpfs/panther/scratch/rob/mongodb-linux-ppc64le-enterprise-rhel71-4.1.4/bin
```

```
if [ ! -d ${LOGDIR} ]
then
  mkdir -p ${LOGDIR}
fi
```

```
if [ ! -d ${DBDIR} ]
then
  mkdir -p ${DBDIR}
fi
```

```
${EXECDIR}/mongod --dbpath ${DBDIR} --logpath ${LOGDIR}/mongo.log --bind_ip_all
```

Lets see it run

```
-bash-4.2$ bsub < mongo_start.lsf  
Job <13556594> is submitted to queue <pantherl>.
```

```
-bash-4.2$ bjobs  
JOBID    USER    STAT  QUEUE      FROM_HOST  EXEC_HOST  JOB_NAME    SUBMIT_TIME  
13556594  rxt62-m  RUN   pantherl  hcplogin2  pgc302@pant  ls-test    May 21 15:34
```

```
-bash-4.2$ bsub -W 00:10 -q pantherl -ls bash
```

Now on the interactive session:

```
bash-4.2$ cd mongodb-linux-ppc64le-enterprise-rhel71-4.1.4/bin/  
./mongo mongodb://pgc302:27017
```

```
MongoDB shell version v4.1.4
```

```
connecting to: mongodb://pgc302:27017
```

```
Implicit session: session { "id" : UUID("5d4b4633-6272-423f-ba2c-def75f835145") }
```

```
MongoDB server version: 4.1.4
```

```
MongoDB Enterprise >
```

The benefits of using MongoDB like this are...

No admin intervention to install

- No requests required
- User can just start work
- Allows user to focus on getting the data

Directly running on the server

- Allows you to use all of the node
- No overheads

Can be shared by teams or even cross teams

- The user who starts the mongodb process is the one writing to the disk
- So usernames can be created in mongodb for members of your team or other teams to write or just read the data

Building Apache with PHP

Building Apache

```
./configure --prefix=/gpfs/panther/scratch/rob/httpd --with-included-apr  
make  
make install
```

Building PHP

```
./configure --prefix=/gpfs/panther/scratch/rob/php --with-apxs2=/gpfs/panther/scratch/rob/httpd/bin/apxs  
make  
make install
```

Alter httpd.conf

```
ServerRoot =/gpfs/panther/scratch/rob/httpd  
DocumentRoot "/gpfs/panther/scratch/rob/httpd/htdocs"  
<Directory "/gpfs/panther/scratch/rob/httpd/htdocs">  
    ScriptAlias /cgi-bin/ "/gpfs/panther/scratch/rob/httpd/cgi-bin/"  
<Directory "/gpfs/panther/scratch/rob/httpd/cgi-bin">
```

Lets start it in a similar way

```
#!/bin/bash
```

```
#BSUB -J ls-test
```

```
#BSUB -oo stdout.txt
```

```
#BSUB -eo stderr.txt
```

```
#BSUB -q panther1
```

```
#BSUB -W 00:05
```

```
#BSUB -n 1
```

```
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/gpfs/panther/scratch/rob/httpd/lib
```

```
export LD_LIBRARY_PATH
```

```
HTTPDIR=/gpfs/panther/scratch/rob/httpd
```

```
${HTTPDIR}/bin/httpd -d ${HTTPDIR} -f ${HTTPDIR}/conf/httpd.conf -k start
```


Summary

Containers are great!

But sometimes it can be better to contain without them

- All resources of a node
- No outside help required
- No need to align with an eco system
- Ability to share work with team or other teams

Acknowledgements

This work was supported by the STFC Hartree Centre Innovation Return on Research programme, funded by the Department for Business, Energy and Industrial Strategy.

Copyright. (c) UKRI–STFC, IBM Corp. 2019

IBM Research @ Hartree Centre

Projects & team

www.research.ibm.com/labs/uk

Location

IBM Research
The Hartree Centre STFC Laboratory
Sci-Tech Daresbury
Warrington WA4 4AD
UK

Contact

Director: martyn.spink@uk.ibm.com
Business enquiries: deborah_sahota@uk.ibm.com



Any Questions?