Pivotal Summit 2019

Bringing Cloud Databases On-Premises with Greenplum and Kubernetes

12–16 November, 2019

Sydney | Singapore | Seoul | Beijing | Tokyo

Part 1: Why

AI and Ease of Use

Pivotal

2011

“Software is eating the world”
Marc Andreessen

2016

“Digital is the main reason just over half of the companies on the Fortune 500 have disappeared since the year 2000”

2019

Software ate the world

AI is ubiquitous

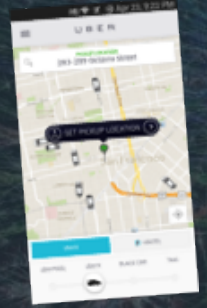

"The importance of accuracy and efficiency […], will continue to rise as we expand and improve products like uberPOOL and beyond."

"On average, customers who have received promotions via Square are more likely to come back and spend 25% more than normal during their next visit."

"Over 75% of what people watch come from our recommendations"

Pivotal

AI is eating software – AI at the core of new startups

BenchSci blog: since November 2017, listed 158 startups using machine learning to research and develop drugs

AI-powered code generation tools like TabNine, TypeSQL and BAYOU

Getting answers to any question about your medical data, from natural language to AI generated SQL (Question-to-SQL)

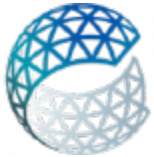

https://www.forbes.com/sites/cognitiveworld/2019/08/29/software-ate-the-world-now-ai-is-eating-software/#388361a85810

Pivotal Customers Bring AI to their core business with Greenplum


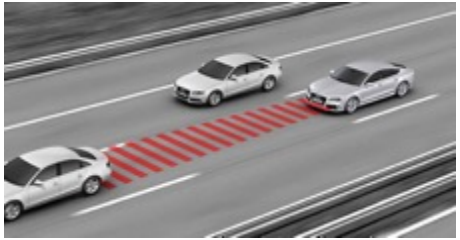









What is AI?

1956: an academic discipline

The “AI Effect”

AI is not just one field

One constant in the definitions is the need for large amount of data, computing power, analytical techniques (e.g. machine learning), an AI platform.


“Every kid coming out of Harvard,…school now thinks he can be the next Mark Zuckerberg, and with these new technologies like cloud computing, he actually has a shot.”

“Having a Shot” at AI with Cloud services

From

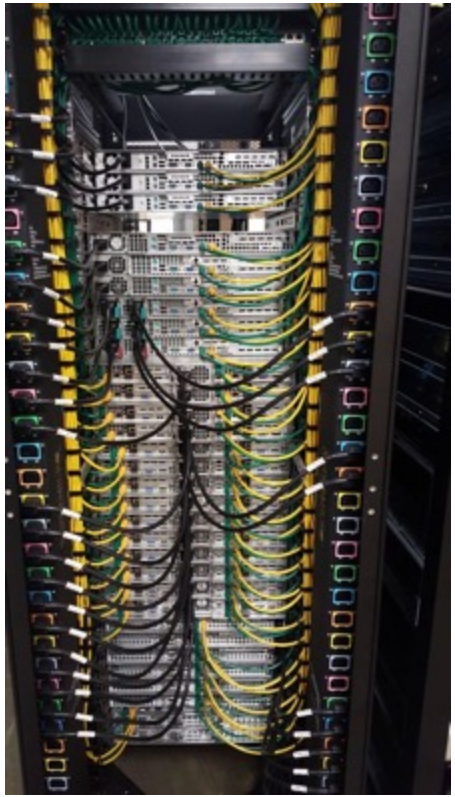

To

But I may not want/be allowed to use the public cloud…


“The cloud is about how you do computing, not where you do computing.”
~ Paul Maritz, Chairman of Pivotal

kubernetes

Container Orchestrator

"Run this containerized app for me. Let me tell you how."

Containerized workloads. Custom and ISV packaged apps and services delivered as containers

Stateful services. Services using persistent storage such as Greenplum, PostgreSQL

Customization. Specify how your app is deployed and operated to optimize performance and reliability

Greenplum for Kubernetes

1. Greenplum is embedded in containers for portability and dependency management
2. Each container is managed by Kubernetes for higher availability & elasticity
3. Kubernetes operator is used for automation
4. PKS for multi-cloud and day-2 operations with full-stack support

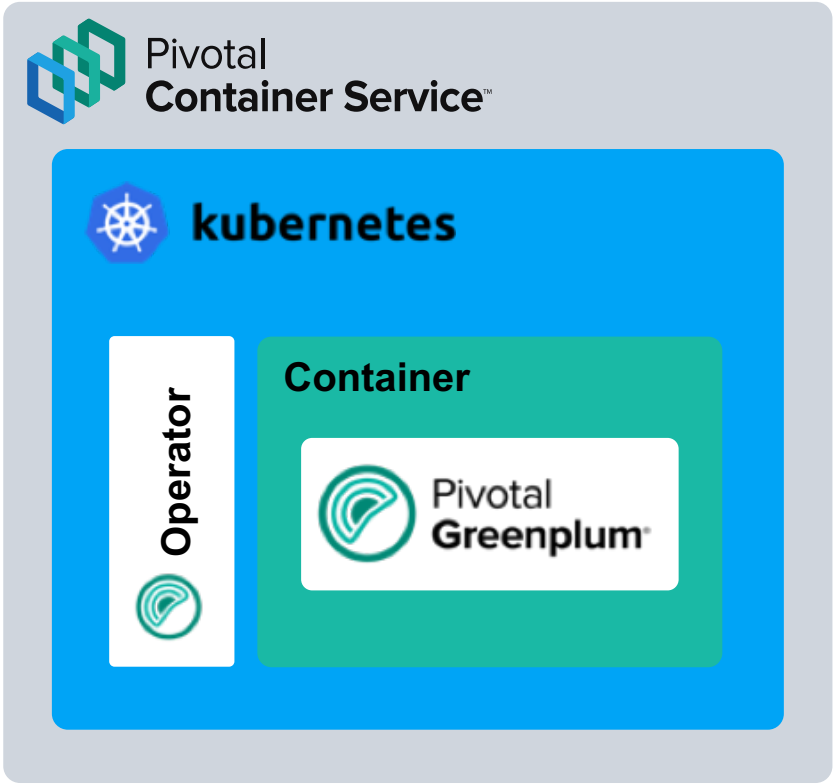

Part 2: Scenario

A day in the Life of a Data Scientist

This is (almost) a real scenario

https://github.com/pnagula/Greenplum-Super-Query


Where is Captain America?

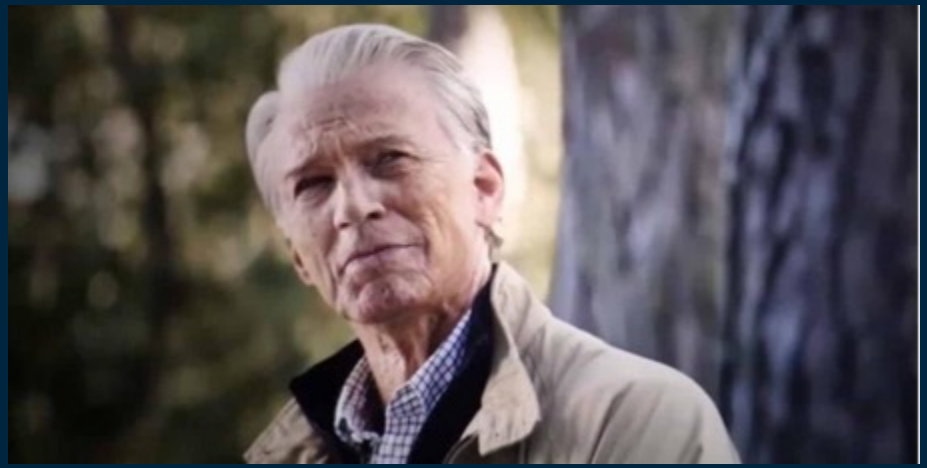


“I can do that all day”

1. I need an AI Platform

Pivotal

My friendly Ops Team has done some “One-Time Setup” for me.

K8s Cluster Ready
Operator Ready

- Downloaded Greenplum for K8s
- Uploaded images to registry
- Created K8s cluster
- Deployed GP Operator
- Prepared instance manifest
  - add extensions
  - adjust storage, sizing, etc.

I get to decide what options to use!

```
apiVersion: "greenplum.pivotal.io/v1"
kind: "GreenplumCluster"
metadata:
  name: my-greenplum
spec:
  masterAndStandby:
    hostBasedAuthentication: |
      # host    all    gpadmin    1.2.3.4/32    trust
      # host    all    gpuser     0.0.0.0/0     md5
    memory: "800Mi"
    cpu: "1"
    storageClassName: standard
    storage: 1G
    antiAffinity: "no"
    workerSelector: {}
  segments:
    primarySegmentCount: 2
    memory: "1800Mi"
    cpu: "1"
    storageClassName: standard
    storage: 1G
    antiAffinity: "no"
    mirrors: "no"
    workerSelector: {}
  gptext:
    serviceName: "my-greenplum-gptext"
  pxf:
    serviceName: "my-greenplum-pxf"
```

For Best Performance:

- Backed by a local SSD
- XFS filesystem, using `readahead` cache

- Only 2 Segments to get started
- 1 GB each because we are in Dev.

.

- No Mirrors we are in Dev.
- AntiAffinity turned off with no mirrors

Same Command

- Initialize Greenplum Workbench
- Update Configuration
- Upgrade Minor Versions
- Apply Patches

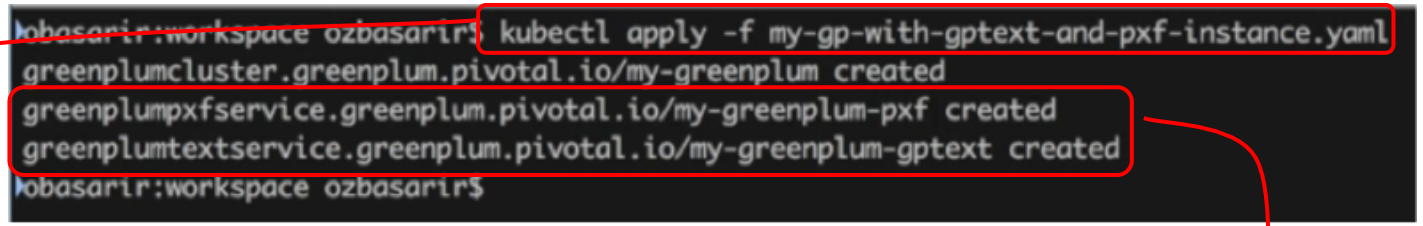

Kubectl apply -f my-gp.yaml

- Options installed automatically

Ready for User Queries in 94 seconds

```
NAME                                                  STATUS    AGE
greenplumcluster.greenplum.pivotal.io/my-greenplum    Running   94s

NAME                                                                AGE
greenplumtextservice.greenplum.pivotal.io/my-greenplum-gptext       93s

NAME                                                              AGE
greenplumpxfservice.greenplum.pivotal.io/my-greenplum-pxf         94s

NAME                                      READY   STATUS    RESTARTS   AGE
pod/greenplum-operator-7fbffdcf64-w6vzw   1/1     Running   0          2d7h
pod/master-0                              1/1     Running   0          90s
pod/master-1                              1/1     Running   0          90s
pod/my-greenplum-gptext-solr-0            1/1     Running   0          93s
pod/my-greenplum-gptext-zookeeper-0       1/1     Running   0          93s
pod/my-greenplum-gptext-zookeeper-1       1/1     Running   0          77s
pod/my-greenplum-gptext-zookeeper-2       1/1     Running   0          53s
pod/my-greenplum-pxf-d5489784b-rhgts      1/1     Running   0          93s
pod/my-greenplum-pxf-d5489784b-sst9n      1/1     Running   0          93s
pod/segment-a-0                           1/1     Running   0          90s
pod/segment-a-1                           1/1     Running   0          90s
```

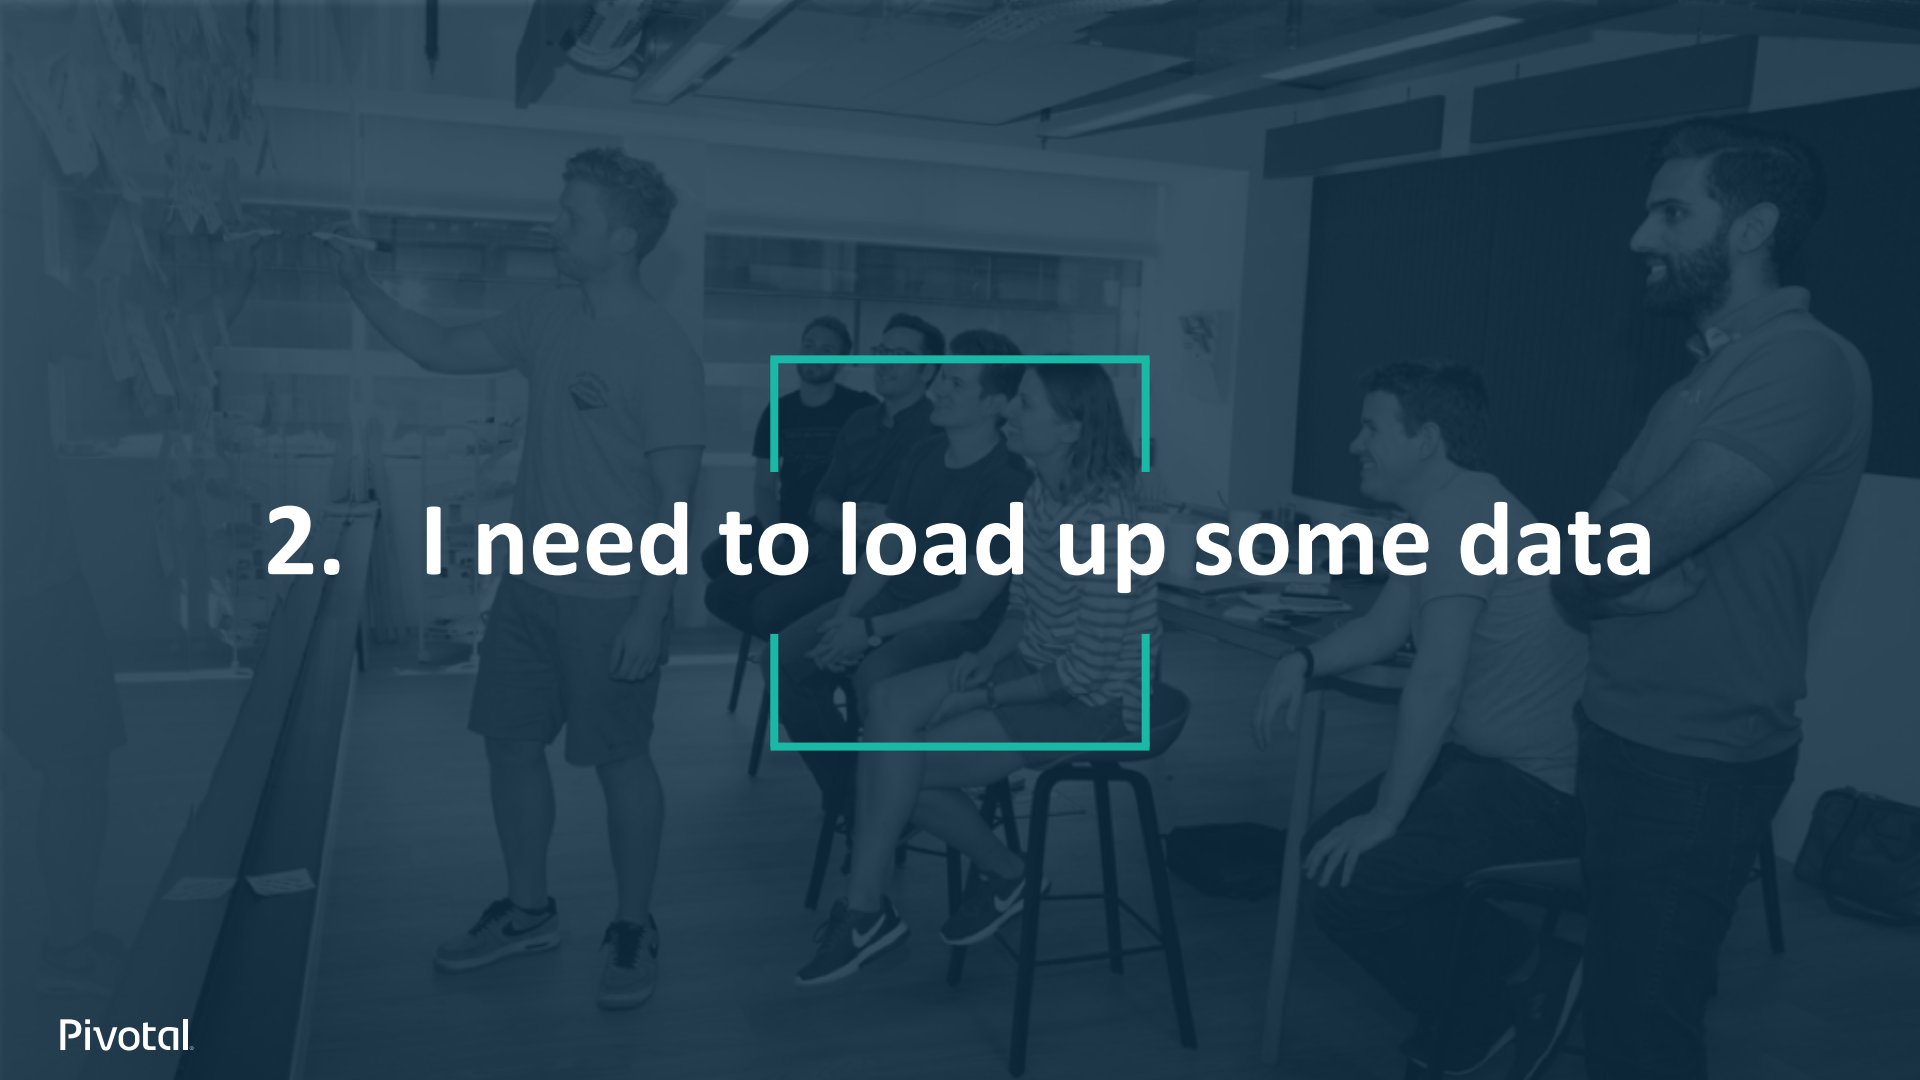

2. I need to load up some data

Pivotal

Greenplum can access it all.

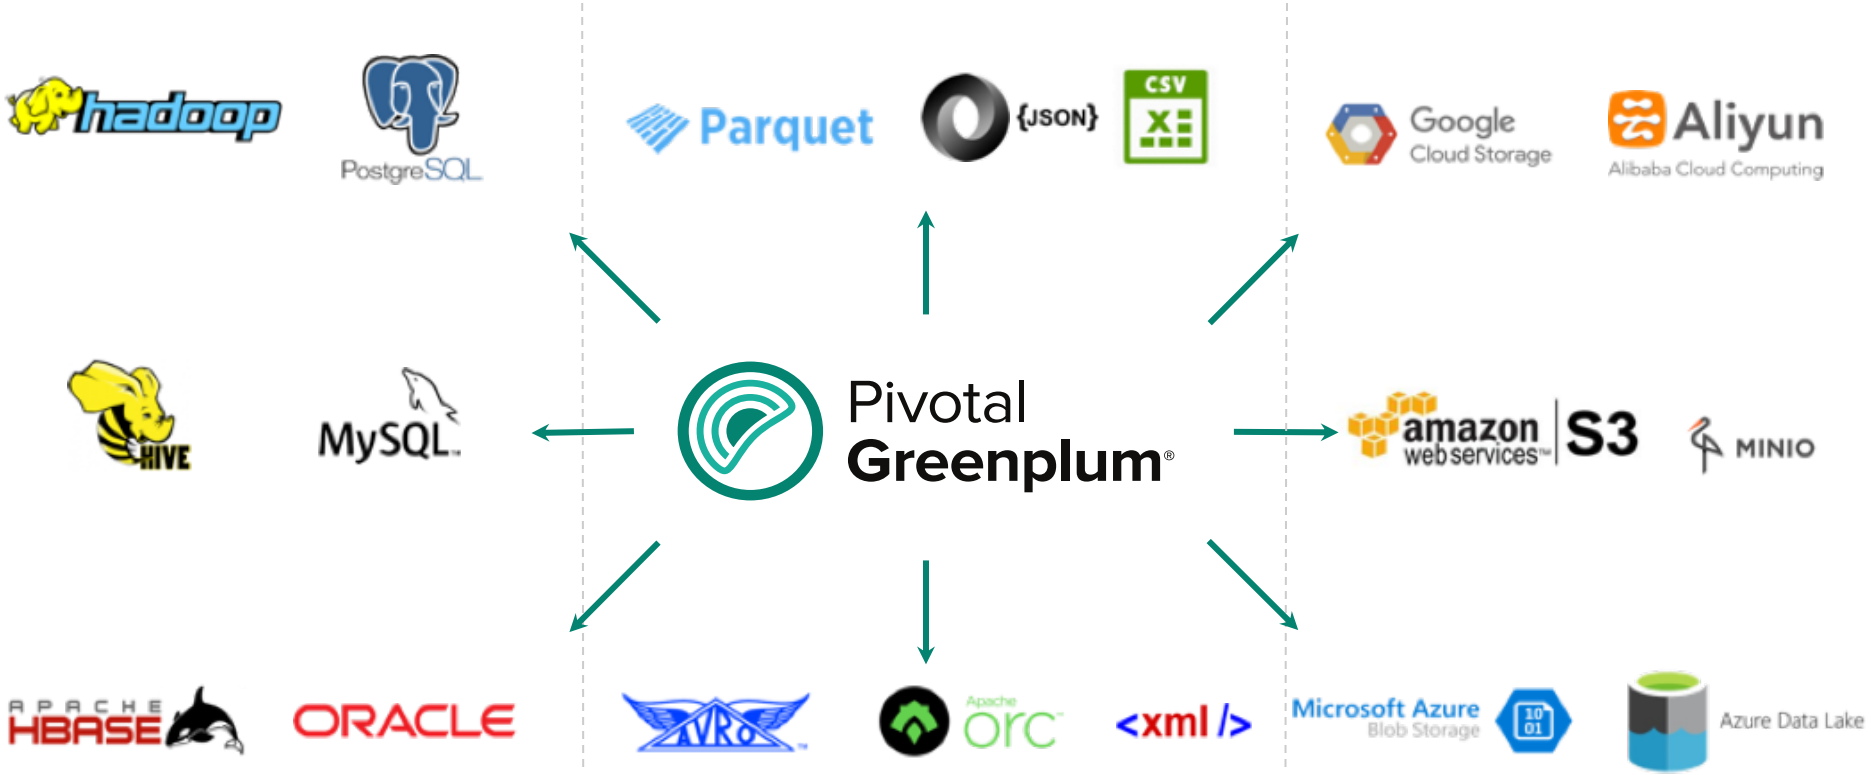

Pivotal Extension Framework (PXF)

- Parallel Access.
- Push Down Processing
- **High Speed (10+ TB/hour) Loading**
- Schema on Read
- Federated Queries
- Standard SQL Interface
- Scale storage independently from compute

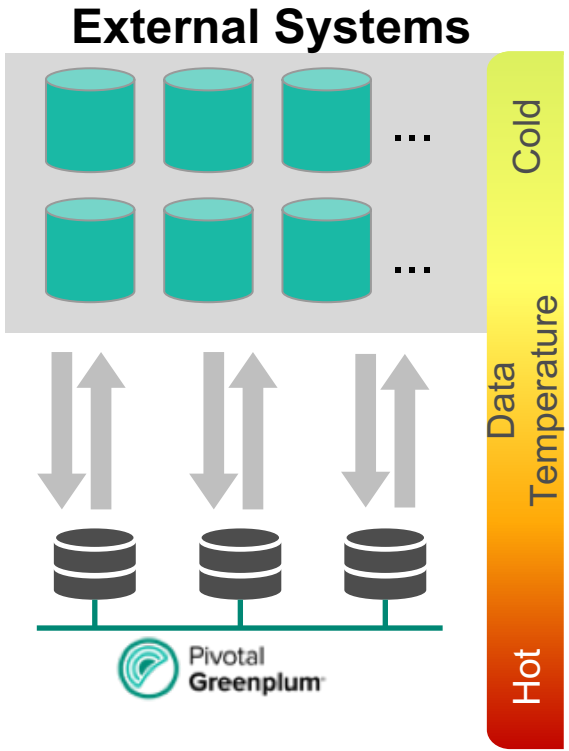

I love my Ops Team - PXF Installed by Default

```
NAME                                              STATUS    AGE
greenplumcluster.greenplum.pivotal.io/my-greenplum   Running   94s

NAME                                                        AGE
greenplumtextservice.greenplum.pivotal.io/my-greenplum-gptext   93s

NAME                                                  AGE
greenplumpxfservice.greenplum.pivotal.io/my-greenplum-pxf   94s

NAME                                        READY   STATUS    RESTARTS   AGE
pod/greenplum-operator-7fbffdcf64-w6vzw     1/1     Running   0          2d7h
pod/master-0                                1/1     Running   0          90s
pod/master-1                                1/1     Running   0          90s
pod/my-greenplum-gptext-solr-0              1/1     Running   0          93s
pod/my-greenplum-gptext-zookeeper-0         1/1     Running   0          93s
pod/my-greenplum-gptext-zookeeper-1         1/1     Running   0          77s
pod/my-greenplum-gptext-zookeeper-2         1/1     Running   0          53s
pod/my-greenplum-pxf-d5489784b-rhgts        1/1     Running   0          93s
pod/my-greenplum-pxf-d5489784b-sst9n        1/1     Running   0          93s
pod/segment-a-0                             1/1     Running   0          90s
pod/segment-a-1                             1/1     Running   0          90s
```

- PXF config is setup automatically
- Scale PXF resources independently of GPDB
- We have installed 2 PXF Servers for HA & Perf.

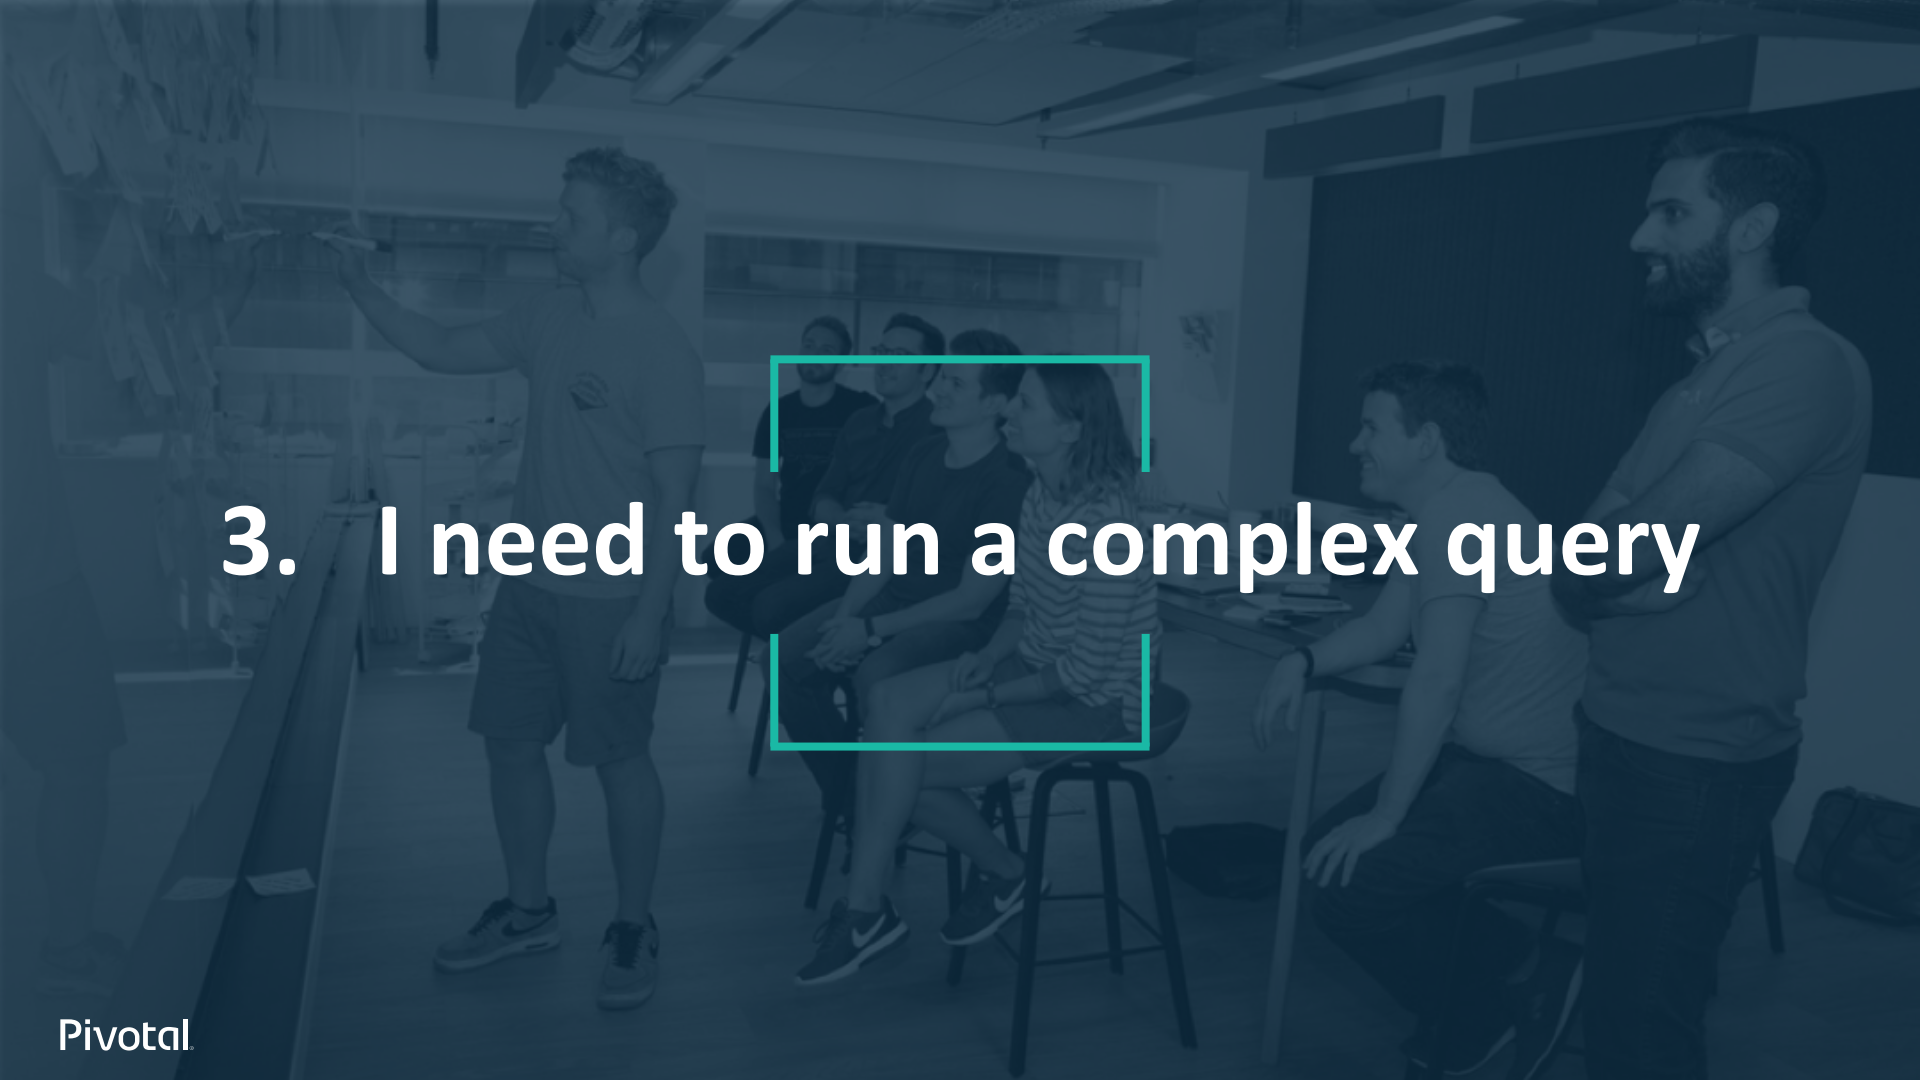

3. I need to run a complex query

Pivotal

User question

“Find anyone whose names sound like ‘Steve’ or “Peggy” and who were at WW2 and knows each other directly and have withdrawn at least $20 after 1945 less than 20 KM from a reference latitude and longitude (Peggy’s parents)”

Unless otherwise indicated, these slides are © 2013-2019 Pivotal Software, Inc. and licensed under a Creative Commons Attribution-NonCommercial license: http://creativecommons.org/licenses/by-nc/3.0/

An interesting Challenge!!!

Language Analytics
Are these the same words?

Graph Analytics
Who do they know?

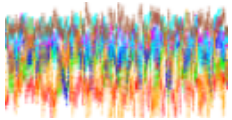

Time Series
When did it happen?

Geospatial Analytics
Where are they?

“Find anyone whose names sound like ‘Steve’ or “Peggy” and who were at WW2 and knows each other directly and have withdrawn at least $20 after 1945 less than 20 KM from a reference latitude and longitude (Peggy’s parents)”

```
drop function if exists get_people(text,text,integer,integer,float,float);
CREATE FUNCTION get_people(text,text,integer,integer,float,float) RETURNS integer
AS $$
declare
linkchk integer; v1 record; v2 record;
begin
 execute 'truncate table results;';
 for v1 in select distinct a.id,a.firstname,a.lastname,amount,tran_date,c.lat,c.lng,address,a.description,d.score from people a,transactions b,location c,
    (SELECT w.id, q.score FROM people w, gptext.search(TABLE(SELECT 1 SCATTER BY 1), 'gpadmin.public.people' , 'World War 2', null) q
     WHERE (q.id::integer) = w.id order by 2 desc) d
     where soundex(firstname)=soundex($1) and a.id=b.id and amount > $3 and (extract(epoch from tran_date) - extract(epoch from now()))/3600 < $4

and st_distance_sphere(st_makepoint($5, $6),st_makepoint(c.lng, c.lat))/1000.0 <= 20.0 and b.locid=c.locid and a.id=d.id
 loop
   for v2 in select distinct a.id,a.firstname,a.lastname,amount,tran_date,c.lat,c.lng,address,a.description,d.score from people a,transactions b,location c,
    (SELECT w.id, q.score FROM people w, gptext.search(TABLE(SELECT 1 SCATTER BY 1), 'gpadmin.public.people' , 'Pivotal', null) q
     WHERE (q.id::integer) = w.id order by 2 desc) d
     where soundex(firstname)=soundex($2) and a.id=b.id and amount > $3 and extract(year from tran_date) > $4
    and st_distance_sphere(st_makepoint($5, $6),st_makepoint(c.lng, c.lat))/1000.0 <= 2.0 and b.locid=c.locid and a.id=d.id
 loop
     execute 'DROP TABLE IF EXISTS out, out_summary;';
     execute  'SELECT madlib.graph_bfs(''people'',''id'',''links'',NULL,'||v1.id||',''out'');'  ;
     select 1 into linkchk from out where dist=1 and id=v2.id;
     if linkchk is not null  then
          insert into results values (v1.id,v1.firstname,v1.lastname,v1.amount,v1.tran_date,v1.lat,v1.lng,v1.address,v1.description,v1.score);
          insert into results values (v2.id,v2.firstname,v2.lastname,v2.amount,v2.tran_date,v2.lat,v2.lng,v2.address,v2.description,v2.score);
      end if;
    end loop;
 end loop;
 return 0;
end
$$ LANGUAGE plpgsql;
--                person1 , person 2, amount, year, longtitude, latitude (in question)
select get_people('Steve','Peggy',20, 1945, 37.926868, -78.024902) ;
```

Greenplum Fuzzy String Match function **Soundex()** to know if people name sounds like ‘Steve’ or ‘Peggy’

GPText.search() function is used to know if both people were at WW2

Greenplum and **Apache MADlib BFS** search to know if there are direct or indirect links between people

Greenplum **Time functions** to calculate amount withdrawn time after the year 1945

Greenplum POSTGIS functions **st_distance_sphere() and st_makepoint()** calculate distance between bank location and reference latitude, longitude < 20 KM

Pivotal

4. I need seamless Day 2 operations

Pivotal

Node Fails - GPDB Auto-recovers

No manual recovery needed;
Just re-run the query!

Master dies and is recovered in 34s.
Same process applies to segments.

Even if its host dies, the master (or segment) will recover on another host because of compute-storage separation.

If you use remote storage then mirrors are not required for auto-recovery.

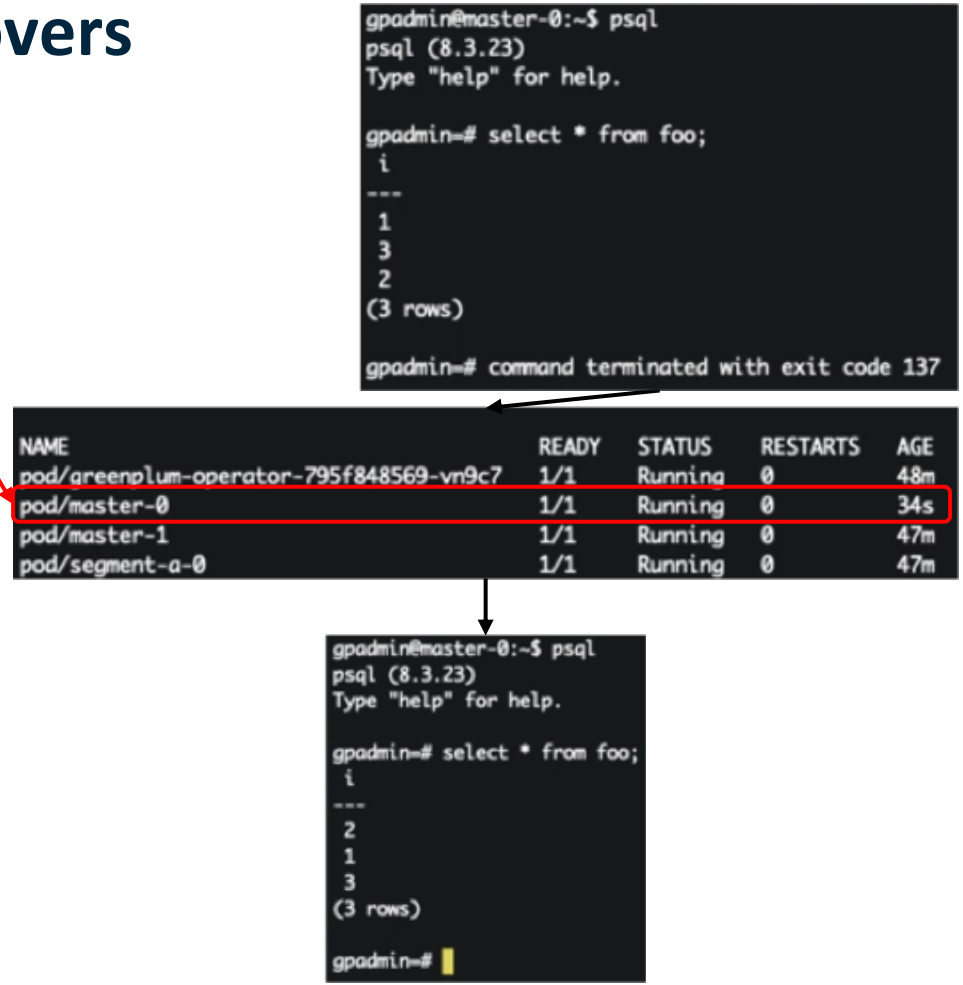

More Data Comes In - Expand GPDB Cluster

- I edit the yaml
- I resubmit the kubectl
- Cluster expands
- GPDB - Autoexpands

- Tested out to 128 Segments
- Linear scaling
- Similar performance to Bare Metal

```
segments:
  primarySegmentCount: 2
```

```
segments:
  primarySegmentCount: 96
```

```
obasarir:workspace ozbasarir$ kubectl apply -f my-gp-with-gptext-and-pxf-instance.yaml
```


I need to optimise cost

Release Compute Resources When Done

Release and Retain State and Data

```
kubectl delete -f my-gp-with-gptext-and-pxf-instance.yaml
```

Patch to a new version

```
kubectl delete -f my-gp-with-gptext-and-pxf-instance.yaml
kubectl apply -f my-gp-with-gptext-and-pxf-instance.yaml
```

Drop Data (Everything gone)

```
kubectl delete pvc --all
```

Part 3: Conclusion

Solution Recap

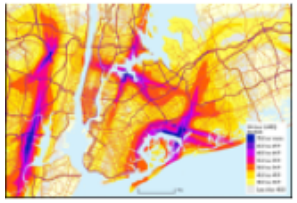

Pivotal
Greenplum®

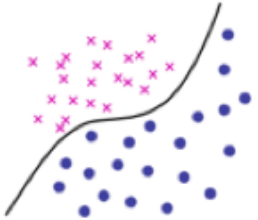

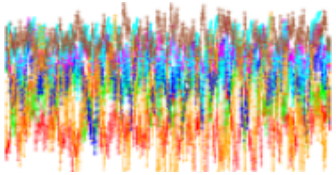

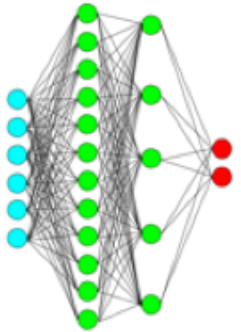

Transforming How The World Builds Software

© Copyright 2019 Pivotal Software, Inc. All rights Reserved.