

SILICON VALLEY



In-Memory
Computing

SUMMIT
2017

DEPLOY LIKE A BOSS: USING APACHE® IGNITE™ AND KUBERNETES®

DANI TRAPHAGEN
SOLUTION ARCHITECT
GRIDGAIN SYSTEMS
@DTRAPEZOID

AGENDA

- Setting up a Apache Ignite cluster
- Using the Kubernetes IP Finder and the Kubernetes Ignite Lookup Service
- Sharing the Ignite Cluster Configuration
- Deploying your Ignite Pods
- Adjusting the Ignite Cluster Size when you need to Scale
 - How to try it out with Azure!

BY THE END – THIS WILL BE YOU WILL BECOME THE...



FIRST - APACHE IGNITE'S PLATFORM

What it isn't:



APACHE IGNITE PLATFORM (VERSION 2.1+ DEPICTED)

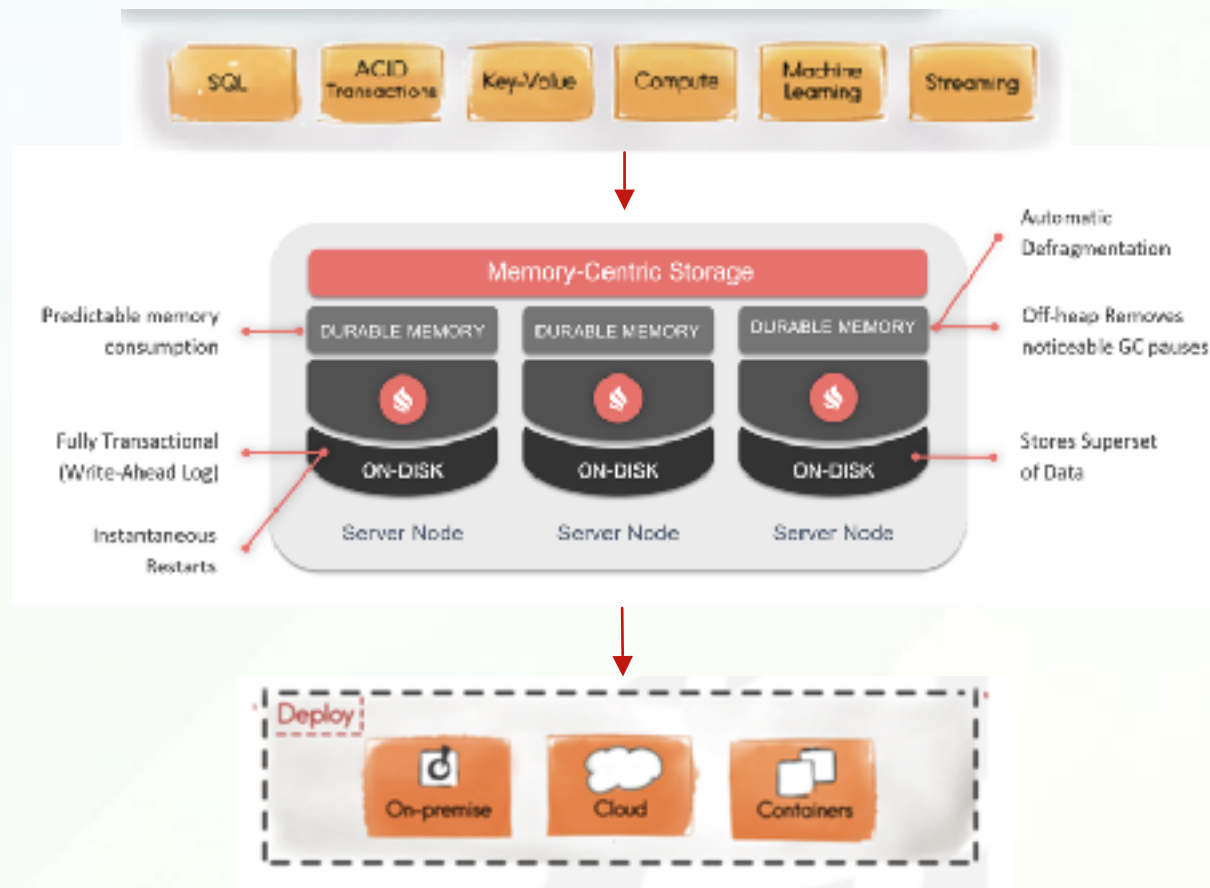
KUBERNETES WORKS WITH VERSIONS 1.9 AND ABOVE

What it is:

APIs to interact with your data

Apache Ignite Cluster

Deployment options




APACHE IGNITE FEATURES

- Ignite Persistence
- ACID Compliance
- Complete SQL Support
- Key-Value Store/APIs
- Collocated Processing
- Scalability and Durability



SECOND – DEPLOYMENT WITH KUBERNETES (K8)

What it isn't:



KUBERNETES IN A SIMPLE DEFINITION

What it is:

- “Kubernetes intends to radically simplify the task of building, deploying and maintaining distributed systems.”
- Kubernetes: Up and Running: Dive into the Future of Infrastructure
 - By: Kelsey Hightower

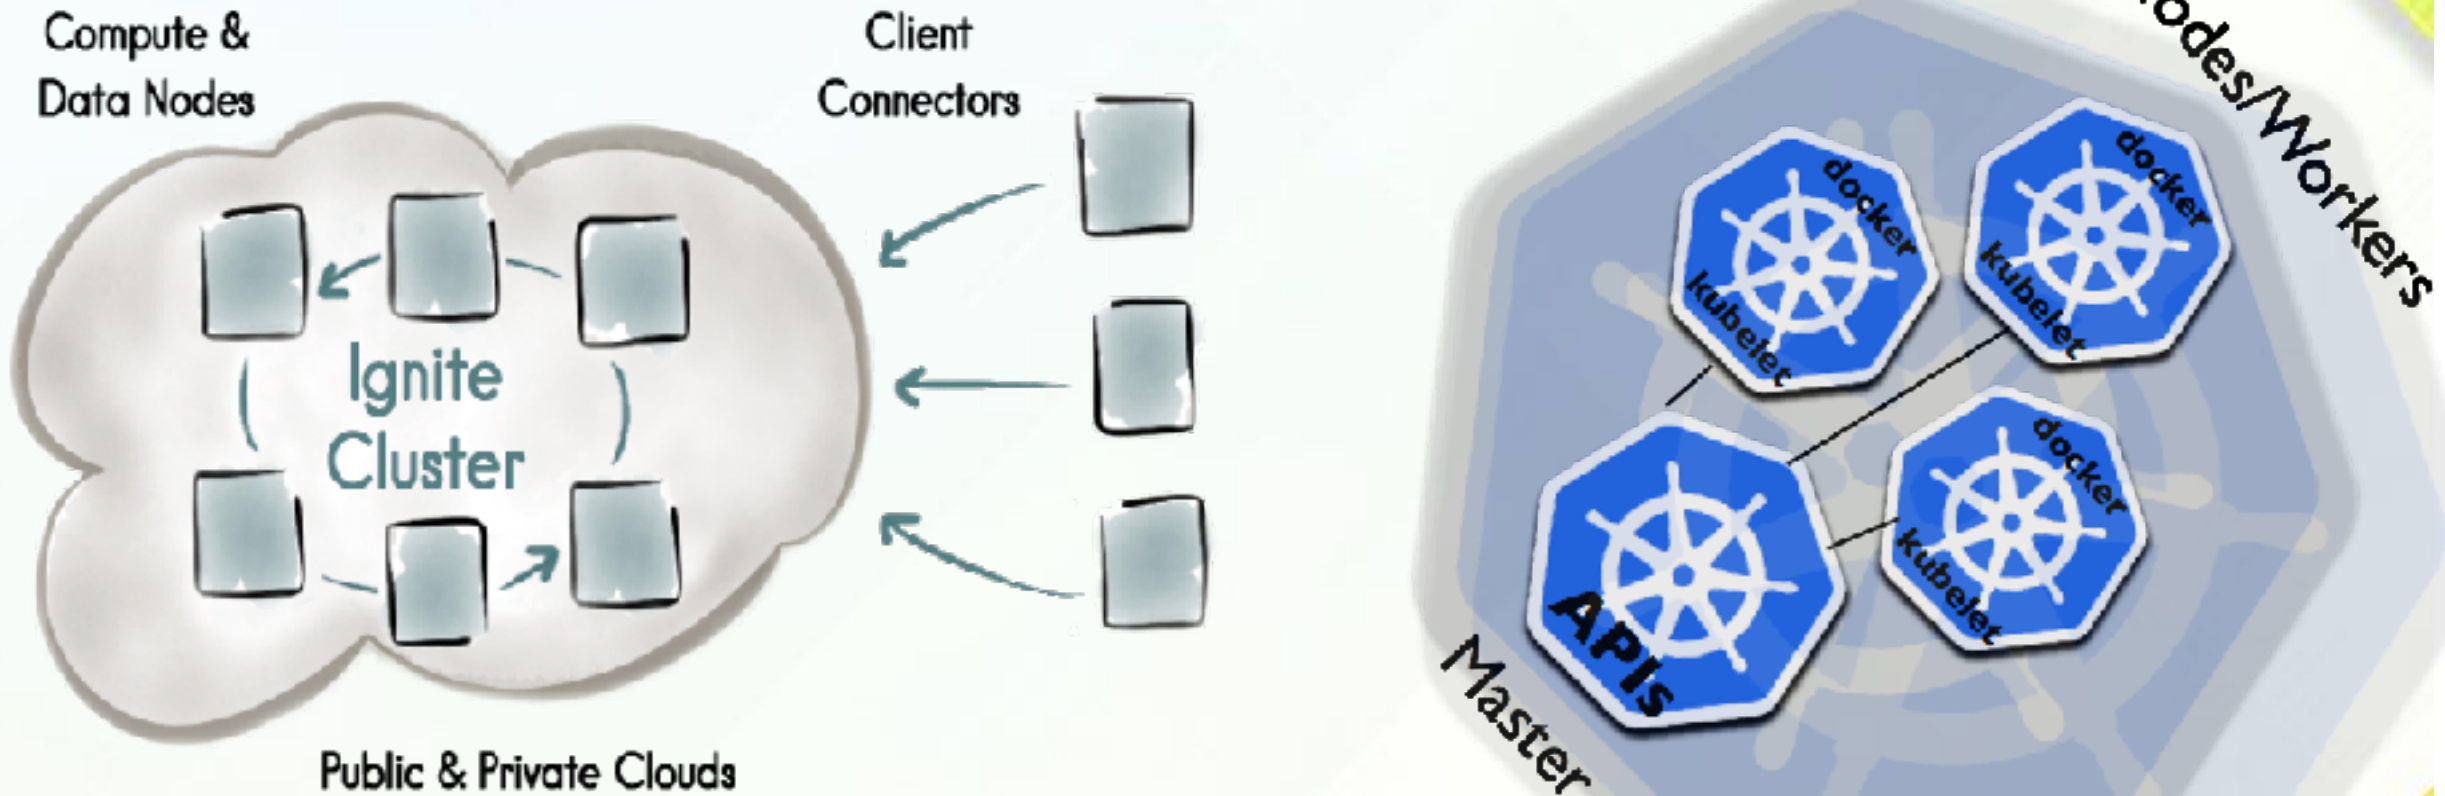


K8 FEATURES

- Automatic Binpacking
- Horizontal Scaling
- Automated Rollouts and Rollbacks
- Storage Orchestration
- Self-Healing
- Service Discovery and Load Balancing
- Secret and Configuration Management
- Batch Execution



APACHE IGNITE + K8 CLUSTER ARCHITECTURE



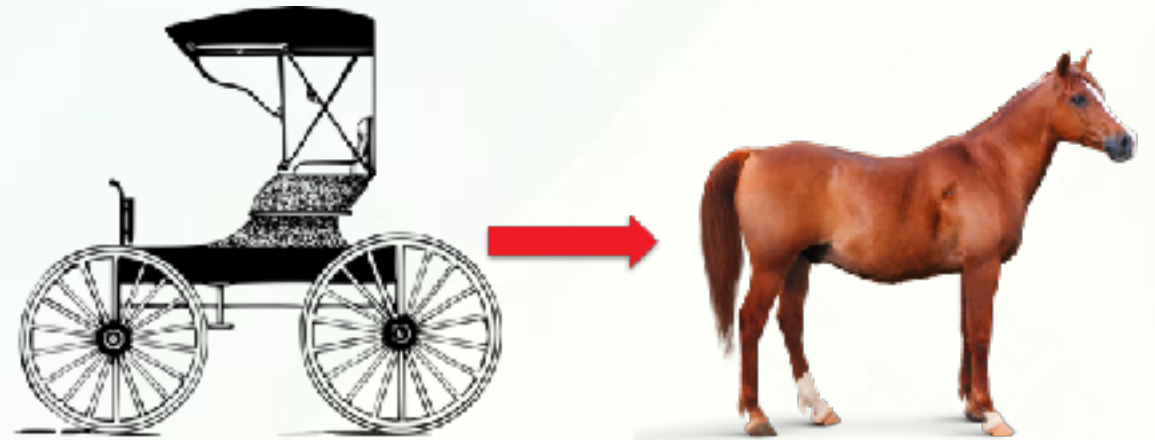
BENEFITS OF K8 (VERSION 1.7)

- Cost Efficiency
 - Use of containers by multiple developers ensures shared resourcing rather than redundancy
- High Availability and Performance
 - *Self-healing*: Deploying with K8 and setting rules for a set number of nodes, will ensure your cluster can always handle the transactions hitting it
 - No more pager duty
 - Get home to your pet chinchilla sooner



SETTING UP AN APACHE IGNITE CLUSTER

- **Step 1:**
- What is your use case?
- Seriously, what is your use case?
- I'm not kidding. Use case, then set up...
 - Cart before horse please.



SETTING UP AN APACHE IGNITE CLUSTER

- 1. [Download](#) Apache Ignite
- 2. Make sure to add the [ignite-kubernetes Maven Dependency](#) to your pom.xml
- 3. Try using examples from [our docs](#).

```
<!--  
    Ignite configuration with all defaults and enabled peer deployment and enabled events.  
-->  
<beans xmlns="http://www.springframework.org/schema/beans"  
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"  
    xsi:schemaLocation="  
        http://www.springframework.org/schema/beans  
        http://www.springframework.org/schema/beans/spring-beans.xsd  
    ">  
  
    <bean class="org.apache.ignite.configuration.IgniteConfiguration">  
        <!-- Enabling the peer-class loading feature. -->  
        <property name="peerClassLoadingEnabled" value="true"/>  
  
        <!--  
            Labeling Data Nodes with special attribute. This attribute is checked by common.filters.DataNodeFilters  
            which decides where caches have to be deployed.  
        -->  
    </bean>  
</beans>
```

SETTING UP K8 ON A LOCAL MACHINE

- Install K8 where you intend
 - AWS
 - Google Cloud
 - Dev Machine
- Set your \$PATH w/K8
 - Install Kubectl
 - [Follow K8 Docs](#)



KUBERNETES DISCOVERY

- Multicast = ☹️
 - Use Static IP Finder & list Ignite IPs, K8 will dynamically assign them
- You can use other cloud Ignite IP finders but you need K8 to running in the cloud env.
- What's the point of the `TcpDiscoveryKubernetesIpFinder`?



USING THE KUBERNETES IP FINDER AND THE KUBERNETES IGNITE LOOKUP SERVICE

```
<property name="discoverySpi">
  <bean class="org.apache.ignite.spi.discovery.tcp.TcpDiscoverySpi">
    <property name="ipFinder">
      <bean class="org.apache.ignite.spi.discovery.tcp.ipfinder.kubernetes.TcpDiscoveryKubernetesIpFinder">
      </bean>
    </property>
  </bean>
</property>
</bean>
```

- Apps & nodes running outside of K8 & Ignite will not be able to reach the cluster
- K8 service should be deployed before Ignite cluster boot
- The Ignite Pods internal IPs will be maintained by the K8 service.
 - Service name must be equal to `setServiceName(String)`
 - This will be `ignite` as a default

USING DOCKER ON LOCAL MACHINE WITH MINIKUBE

```
danitrophagen@Dani-NBP ~$ docker ps
```

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS
97cc92b952b8	gcr.io/google_containers/kubernetes-dashboard-amd64	"/dashboard --inse..."	3 minutes ago	Up 3 minutes
abeb3cac70c2	gcr.io/google_containers/pause-amd64:3.0	"/pause"	4 minutes ago	Up 4 minutes
6a0a5bd1fe67	gcr.io/google_containers/k8s-dns-dnsmasq-nanny-amd64	"/dnsmasq-nanny -v..."	35 minutes ago	Up 35 minutes
61c0b0ec07d4	gcr.io/google_containers/k8s-dns-kube-dns-amd64	"/kube-dns --comai..."	35 minutes ago	Up 35 minutes
02282e7c7c22	gcr.io/google_containers/k8s-dns-sidecar-amd64	"/sidecar --v=2 --..."	35 minutes ago	Up 35 minutes
c4c5a18dcd4d	gcr.io/google_containers/pause-amd64:3.0	"/pause"	35 minutes ago	Up 35 minutes
b6743ee53b24	gcr.io/google_containers/echoserver	"nginx -g 'daemon ...'"	About an hour ago	Up 47 minutes
d8c59a7c5325	gcr.io/google_containers/pause-amd64:3.0	"/pause"	About an hour ago	Up 47 minutes
Fa27860f8345	gcr.io/google_containers/kube-addon-manager	"/opt/kube-addons.sh"	About an hour ago	Up 55 minutes
157a056988b7	gcr.io/google_containers/pause-amd64:3.0	"/pause"	About an hour ago	Up 55 minutes

```
danitrophagen@Dani-NBP ~$
```


SHARING THE IGNITE CLUSTER CONFIGURATION

Sharing is caring



SERVICE STARTUP & SHARING CONFIGS

- Minikube - Two basic commands
 - minikube start
 - minikube dashboard
- Kubectl
 - kubectl create -f ~/<path-to-project>/<project-name>/config/ignite-service.yaml

```
danitraphagen@Dani-MBP ~$ minikube start
Starting local Kubernetes v1.7.0 cluster...
Starting VM...
Getting VM IP address...
Moving files into cluster...
Setting up certs...
Starting cluster components...
Connecting to cluster...
Setting up kubeconfig...
Kubectl is now configured to use the cluster.
danitraphagen@Dani-MBP ~$ minikube dashboard
Opening kubernetes dashboard in default browser...
danitraphagen@Dani-MBP ~$
```

```
service "ignite" created
danitraphagen@Dani-MBP ~$
```



Yes, someone I follow on Twitter did this.

NOW KUBERNETES IS RUNNING!

The screenshot shows the Kubernetes dashboard interface. At the top, there is a search bar and a '+ CREATE' button. The left sidebar contains a navigation menu with categories like Cluster, Namespaces, Nodes, Persistent Volumes, Roles, Storage Classes, and Workloads. The 'Workloads' section is expanded, showing three sub-sections: Deployments, Pods, and Replica Sets. Each section contains a table with columns for Name, Labels, Pods, Age, and Images. The 'hello-minikube' deployment is shown with 1/1 pods, and the 'hello-minikube-130744149' pod is shown in a 'Running' state.

Deployments

Name	Labels	Pods	Age	Images
hello-minikube	run: hello-minikube	1 / 1	57 minutes	gcr.io/google_containers/echoserver:1.4

Pods

Name	Status	Restarts	Age
hello-minikube-130744149-2rtg5	Running	1	57 minutes

Replica Sets

Name	Labels	Pods	Age	Images
hello-minikube-130744149	pod-template-hash: 130744149 run: hello-minikube	1 / 1	57 minutes	gcr.io/google_containers/echoserver:1.4

CONFIGURING YOUR IGNITE PODS

■ 2 Things Needed!

- 1. [Apache Ignite Configuration File](#) with the Kubernetes IP Finder
- 2. [YAML Configurations](#) for the Apache Ignite pods/nodes

■ Steps:

■ Create your ignite-service.yaml

- `kubectl create -f ignite-service.yaml`
- `kubectl get svc ignite`



NAME	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
ignite	None	<none>	9042/TCP	29s

SHARING IGNITE CLUSTER CONFIGS

- Confirm the ignite service was created
- Make a path to the persistence volume docker will use to pass the kubernetes config 'example-kube.xml'

```
danitraphagen@Dani-MBP ~$ kubectl get svc ignite
```

NAME	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
ignite	None	<none>	9042/TCP	13m

```
danitraphagen@Dani-MBP ~/data/ignite$ kubectl create -f ignite-volume.yaml  
persistentvolume "ignite-volume" created
```

SUCCESS

Now, do this for the
ignite-volume-claim.yaml



```
danitraphagen@Dani-MBP ~/data/ignite$ kubectl get pv ignite-volume  


| NAME          | CAPACITY | ACCESSMODES | RECLAIMPOLICY | STATUS    |
|---------------|----------|-------------|---------------|-----------|
| ignite-volume | 1Gi      | RWO         | Retain        | Available |

  
danitraphagen@Dani-MBP ~/data/ignite$
```

ignite-volume-claim.yaml

```
kind: PersistentVolumeClaim  
apiVersion: v1  
metadata:  
  name: ignite-volume-claim  
spec:  
  accessModes:  
    - ReadWriteOnce  
  resources:  
    requests:  
      storage: 1Gi
```


PERSISTENT VOLUME BOUND?

- Make sure your persistent volume is bound to the claim
 - `kubectl get pvc ignite-volume-claim`
 - `kubectl get pv ignite-volume`

NAME	STATUS	VOLUME
ignite-volume-claim	Bound	ignite-volume

STATUS	CLAIM
Bound	default/ignite-volume-claim

An arrow points from the 'Bound' status in the first table to the 'Bound' status in the second table.



use xhyve driver so you can access persistent volumes

and include start on my computer="/data/ignite" that is used in the examples to **"/Users/<username>/data/ignite"**

DEPLOYING YOUR IGNITE PODS



- Now it's time to launch
 - `kubectl create -f ignite-deployment.yaml`
 - `kubectl get pods`
- Get the logs and examine for each cluster...
ex)
 - `kubectl logs ignite-cluster-3454482164-d4m6g`
- Scale out:
 - `kubectl scale --replicas=5 -f ignite-deployment.yaml`

NAME	READY	STATUS	RESTARTS	AGE
ignite-cluster-3454482164-d4m6g	1/1	Running	0	25m
ignite-cluster-3454482164-w0xtx	1/1	Running	0	25m

NAME	READY	STATUS	RESTARTS	AGE
ignite-cluster-3454482164-d4m6g	1/1	Running	0	34m
ignite-cluster-3454482164-ktkrr	1/1	Running	0	58s
ignite-cluster-3454482164-r20f8	1/1	Running	0	58s
ignite-cluster-3454482164-vf8kh	1/1	Running	0	58s
ignite-cluster-3454482164-w0xtx	1/1	Running	0	34m

MANY DEPLOYMENT OPTIONS

- On Premise
- Cloud
 - Azure
 - EC2
 - Google Cloud

```
Last login: Wed Aug 23 01:30:22 2017 from cpe-172-114-236-18.socal.res.nn.com

  _| _|_ )
  _| (   /  Amazon Linux AMI
  ___|\___|___|

https://aws.amazon.com/amazon-linux-ami/2017.03-release-notes/
[ec2-user@ip-172-31-12-255 ~]$ curl -LO https://storage.googleapis.com/kubernetes-release/release/$(curl -s https://storage.googleapis.com/kubernetes-release/release/stable.txt)/bin/linux/amd64/kubectl
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload  Total   Spent    Left  Speed
100 68.9M  100 68.9M    0     0  36.0M      0  0:00:01  0:00:01 --:--:-- 36.0M
[ec2-user@ip-172-31-12-255 ~]$ chmod +x ./kubectl
[ec2-user@ip-172-31-12-255 ~]$ sudo mv ./kubectl /usr/local/bin/kubectl
[ec2-user@ip-172-31-12-255 ~]$ kubectl cluster-info
Kubernetes master is running at http://localhost:8080

To further debug and diagnose cluster problems, use 'kubectl cluster-info dump'.
[ec2-user@ip-172-31-12-255 ~]$
```


USING MINIKUBE FOR LOCAL DEV

- A good place to start for exploration
- When you want to get in the cloud - pick your poison
 - Post on setup w/Azure by Ignite PMC Denis Magda:
 - <https://dzone.com/articles/deploying-apache-ignite-in-kubernetes-on-microsoft>
 - <https://kubernetes.io/docs/setup/pick-right-solution/#turnkey-cloud-solutions>

```
danitrophager@Dani-MBP ~$ curl -Lo minikube https://storage.googleapis.com/minikube/releases/v0.21.0/minikube-darwin-amd64 && chmod +x minikube && sudo mv minikube /usr/local/bin/

  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100 81.9M  100 81.9M    0     0 8189k      0  0:00:10  0:00:10 --:--:-- 8865k
danitrophager@Dani-MBP ~$ minikube start
Starting local Kubernetes v1.7.0 cluster...
Starting VM...
Downloading Minikube ISO
 97.80 MB / 97.80 MB [=====] 100.00% 0s
Getting VM IP address...
Moving files into cluster...
Setting up certs...
Starting cluster components...
Connecting to cluster...
Setting up kubeconfig...
Kubectl is now configured to use the cluster.
```

DARE TO TRY?



1. DEPLOY CLOUD ENVIRONMENT

```
dani@Azure:~$ kubectl get nodes
NAME                                STATUS              AGE           VERSION
k8s-agent-747b6f74-0               NotReady            15s           v1.6.6
k8s-agent-747b6f74-1               NotReady            15s           v1.6.6
k8s-agent-747b6f74-2               NotReady            19s           v1.6.6
k8s-master-747b6f74-0              NotReady,SchedulingDisabled 17s           v1.6.6
dani@Azure:~$ kubectl get nodes
NAME                                STATUS              AGE           VERSION
k8s-agent-747b6f74-0               Ready                1m           v1.6.6
k8s-agent-747b6f74-1               Ready                1m           v1.6.6
k8s-agent-747b6f74-2               Ready                1m           v1.6.6
k8s-master-747b6f74-0              Ready,SchedulingDisabled 1m           v1.6.6
dani@Azure:~$ █
```


2. CONNECT TO YOUR CLOUD ENVIRONMENT FROM YOUR LOCAL MACHINE, EX) AZURE

#connect to cluster, make sure your sshkeys are setup

```
az acs kubernetes get-credentials --resource-  
group=myResourceGroup --name=myK8sCluster
```

#make sure you see the k8s-agents & master

```
kubectl get nodes
```

USING THE DASHBOARD

- 1.kubectl proxy
- 2.http://localhost:8001/ui

The screenshot shows the Kubernetes dashboard interface. At the top, there is a blue header with the 'kubernetes' logo and a breadcrumb trail 'Admin > Nodes'. On the left side, there is a navigation menu with sections: 'Admin' (containing 'Namespaces', 'Nodes', and 'Persistent Volumes'), 'Namespace' (set to 'default'), 'Workloads' (containing 'Deployments', 'Replica Sets', 'Replication Controllers', 'Daemon Sets', 'Stateful Sets', 'Jobs', and 'Pods'), and 'Services and discovery'. The main content area is divided into two columns. The left column contains a 'CPU usage' line graph showing CPU scores over time from 13:05 to 13:22. The right column contains a 'Mem' (Memory) section with a vertical scale from 3.26 to 14.7 bytes. Below these, there is a 'Nodes' section with a table listing node names and their labels. One node, 'k8s-agent-2ed260dc-0', is highlighted with a green checkmark. The labels for this node include 'beta.kubernetes.io/arch', 'beta.kubernetes.io/instanceid', 'beta.kubernetes.io/os', and 'failure-domain.beta.kubernetes.io/zone'. A 'show all labels' link is provided at the bottom of the node list.

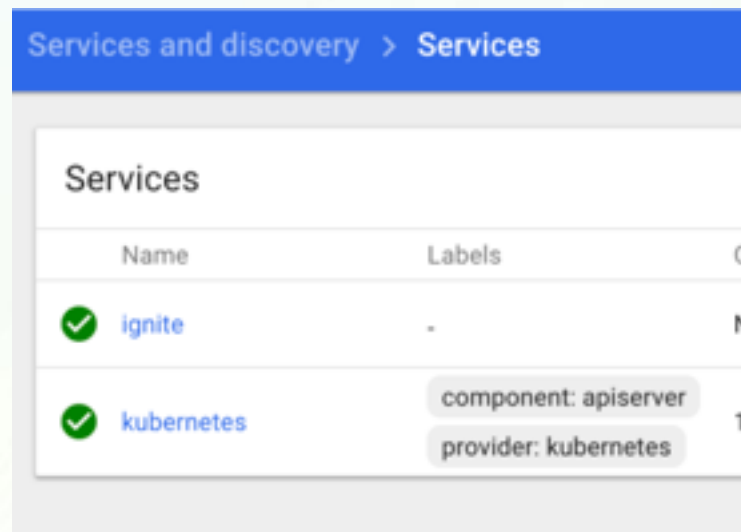
Name	Labels
<input checked="" type="checkbox"/> k8s-agent-2ed260dc-0	<ul style="list-style-type: none">beta.kubernetes.io/archbeta.kubernetes.io/in...beta.kubernetes.io/os...failure-domain.beta.k...failure-domain.beta.k... show all labels

4. CREATE THE K8 LOOKUP SERVICE

#using above link, next to the number 4, create the file then initiate the service

```
kubectl create -f ignite-service.yaml
```

#you will see that the ignite service is under the services tab



Services and discovery > Services

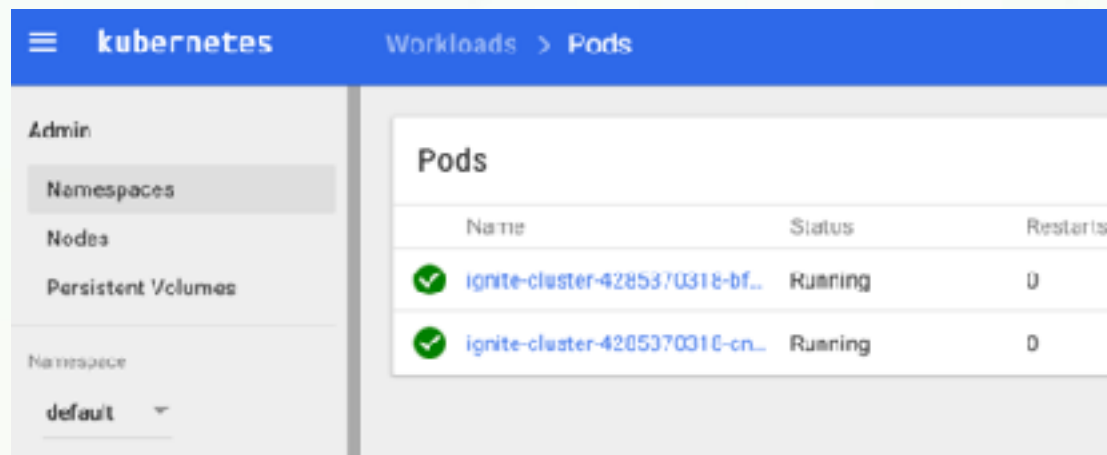
Services		
Name	Labels	C
✓ ignite	-	N
✓ kubernetes	component: apiserver provider: kubernetes	1

5. DEPLOY YOUR APACHE IGNITE CLUSTER

#create the ignite-deployment.yaml file following instructions [here](#)

```
kubectl create -f ignite-deployment.yaml
```

#you will see that the ignite cluster is running in kubernetes



Name	Status	Restarts
ignite-cluster-4285370318-bf...	Running	0
ignite-cluster-4285370310-cn...	Running	0

ADJUSTING THE IGNITE CLUSTER SIZE WHEN YOU NEED TO SCALE

- When you want to elastically scale out your cluster with K8:

- `kubectl scale --replicas=<n> -f ignite-deployment.yaml`

- Run:

- `kubectl get pods`

- Let's say you want 5 nodes?

- `kubectl scale --replicas=5 -f ~/kubernetes_dev/azure/ignite-deployment.yaml`

- You will see your cluster scale out – in this case from 2 to 5 nodes!



OVERALL STEPS

Ignite Download:

<https://ignite.apache.org/download.cgi>

Run Kubernetes Locally:

<https://kubernetes.io/docs/getting-started-guides/minikube/>

Deploy Kubernetes & Ignite

<https://apacheignite.readme.io/docs/kubernetes-deployment>

RESOURCES

- Denis Magda's & Akmal Chaudhri's blog's on deploying Ignite w/K8 and Azure or AWS:
 - <https://dzone.com/articles/deploying-apache-ignite-in-kubernetes-on-microsoft>
 - <https://www.gridgain.com/resources/blog/kubernetes-and-apacher-ignitetm-deployment-aws>
- Tutorials
 - <https://kubernetes.io/docs/tutorials/>
- K8 Book by Kelsey Hightower:
 - <http://shop.oreilly.com/product/0636920043874.do>
- Ignite Book by Shahim & others:
 - <https://leanpub.com/ignite>



ANY QUESTIONS?



- Reach me on Twitter @dtrapezoid
- Use #apacheignite
- Checkout <http://ignite.apache.org>

