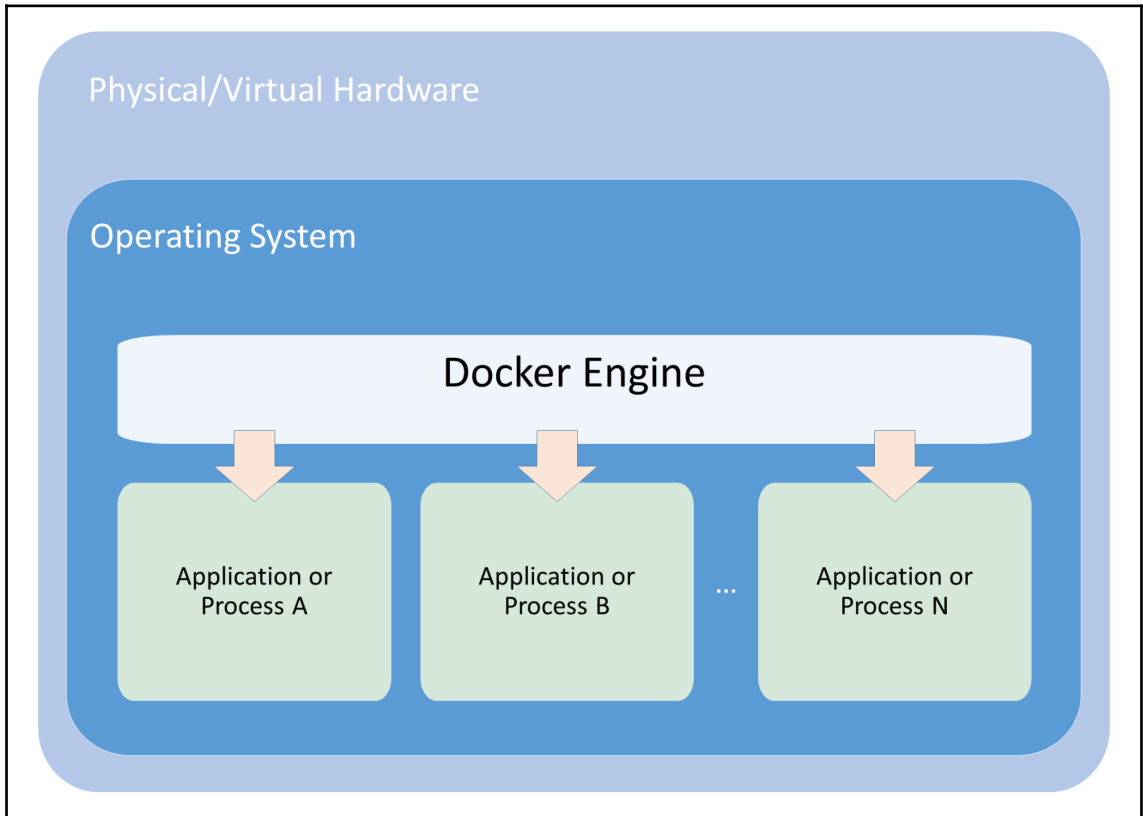
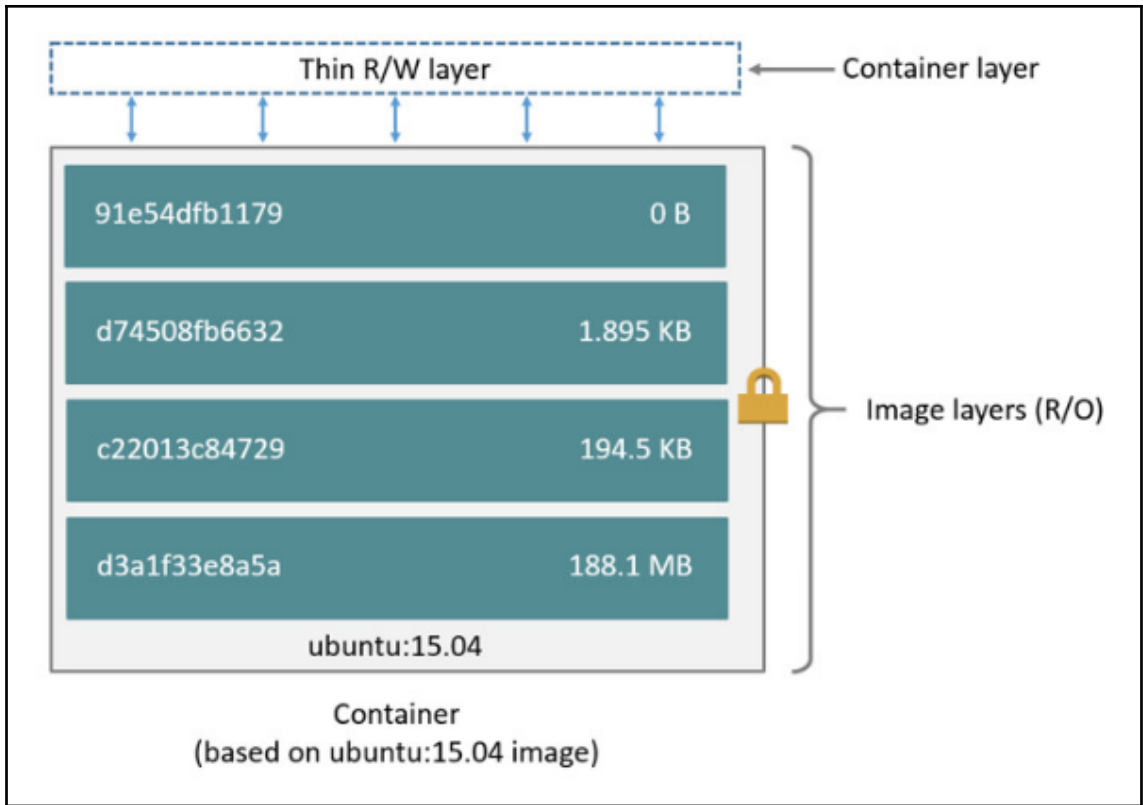
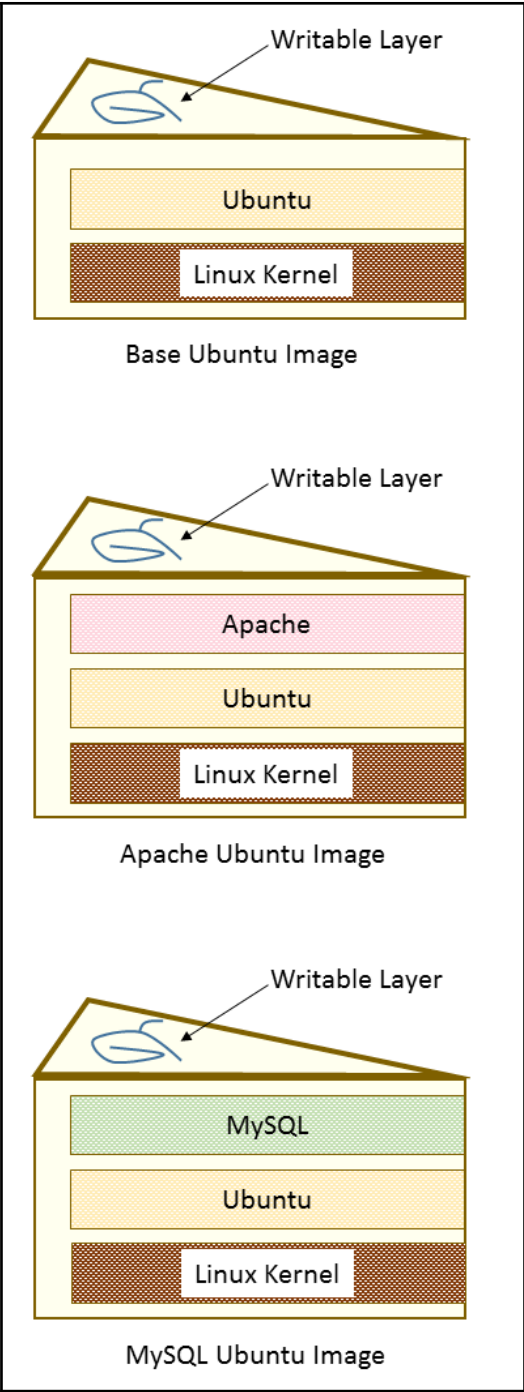
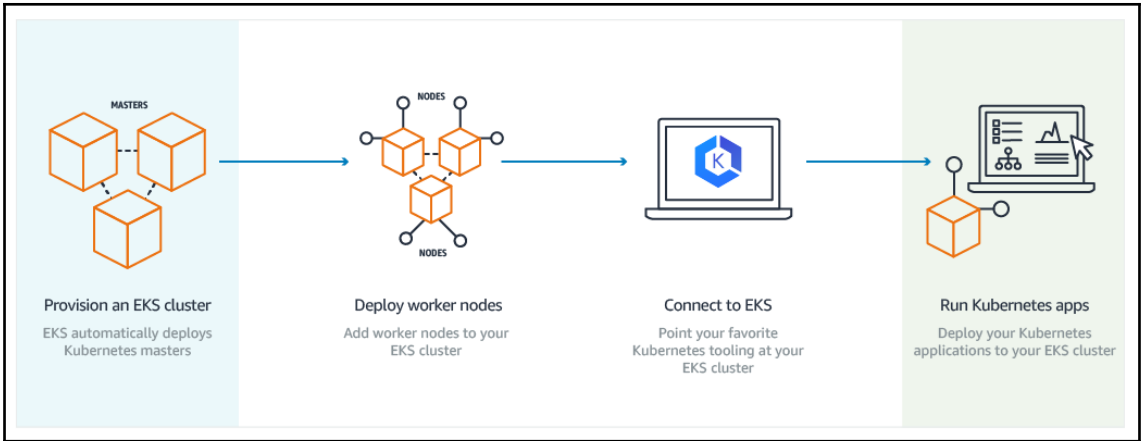


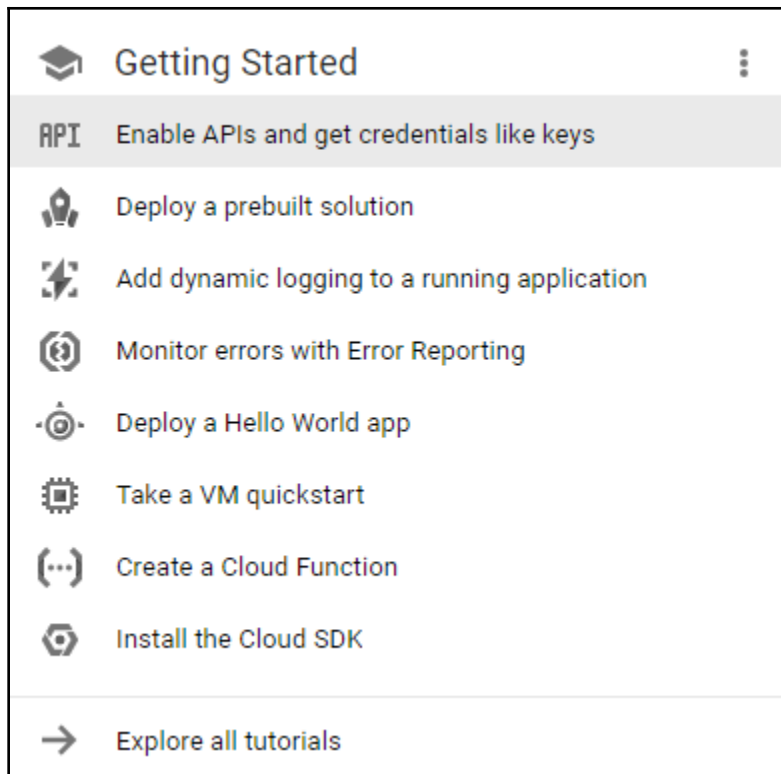
Chapter 1: Introduction to Kubernetes











```
... Starting cluster in us-central1-b using provider gce  
... calling verify-prereqs
```

```
All components are up to date.
```

```
All components are up to date.
```

```
All components are up to date.
```

```
... calling kube-up
Your active configuration is: [default]

Project: dynamic-nomad-152102
Zone: us-central1-b
gs://kubernetes-staging-549d6b8d9c/kubernetes-devel/
+++ Staging server tars to Google Storage: gs://kubernetes-staging-549d6b8d9c/kubernetes-devel
+++ kubernetes-server-linux-amd64.tar.gz uploaded (sha1 = 5df19e3745bbc8c7d1a5bf6d61d9e1b0d189db64)
+++ kubernetes-salt.tar.gz uploaded (sha1 = 95e855d893e4549b935aed8736f3a2372ae7cd3)
+++ kubernetes-manifests.tar.gz uploaded (sha1 = e9c52530a14612c91f45e017743925a0dba6dcc8)
INSTANCE_GROUPS=
NODE_NAMES=
```

```
Looking for already existing resources
Starting master and configuring firewalls
Created [https://www.googleapis.com/compute/v1/projects/dynamic-nomad-152102/zones/us-central1-b/disks/kubernetes-master-pd].
```

NAME	ZONE	SIZE_GB	TYPE	STATUS
kubernetes-master-pd	us-central1-b	20	pd-ssd	READY

New disks are unformatted. You must format and mount a disk before it can be used. You can find instructions on how to do this at:

<https://cloud.google.com/compute/docs/disks/add-persistent-disk#formatting>

```
Created [https://www.googleapis.com/compute/v1/projects/dynamic-nomad-152102/global/firewalls/kubernetes-master-https].
```

NAME	NETWORK	SRC_RANGES	RULES	SRC_TAGS	TARGET_TAGS
kubernetes-master-https	default	0.0.0.0/0	tcp:443		kubernetes-master

```
Created [https://www.googleapis.com/compute/v1/projects/dynamic-nomad-152102/regions/us-central1/addresses/kubernetes-master-ip].
```

```
Generating certs for alternate-names: IP:23.251.158.223,IP:10.0.0.1,DNS:kubernetes,DNS:kubernetes.default,DNS:kubernetes.default.svc,DNS:kubernetes.default.svc.cluster.local,DNS:kubernetes-master
```

```

+++ Logging using Fluentd to gcp
WARNING: You have selected a disk size of under [200GB]. This may result in poor
I/O performance. For more information, see: https://developers.google.com/compute/docs/disks#pdperformance.
Created [https://www.googleapis.com/compute/v1/projects/dynamic-nomad-152102/global/firewalls/kubernetes-minion-all].
NAME                NETWORK  SRC_RANGES      RULES                               SRC_TAGS
TARGET_TAGS
kubernetes-minion-all  default  10.244.0.0/14   tcp,udp,icmp,esp,ah,sctp
kubernetes-minion
Created [https://www.googleapis.com/compute/v1/projects/dynamic-nomad-152102/zones/us-central1-b/instances/kubernetes-master].
NAME                ZONE     MACHINE_TYPE    PREEMPTIBLE  INTERNAL_IP  EXTERNAL_IP  STATUS
kubernetes-master  us-central1-b  n1-standard-1          10.128.0.2   23.25
1.158.223  RUNNING
Creating minions.
Attempt 1 to create kubernetes-minion-template
WARNING: You have selected a disk size of under [200GB]. This may result in poor
I/O performance. For more information, see: https://developers.google.com/compute/docs/disks#pdperformance.
Created [https://www.googleapis.com/compute/v1/projects/dynamic-nomad-152102/global/instanceTemplates/kubernetes-minion-template].
NAME                MACHINE_TYPE    PREEMPTIBLE  CREATION_TIMESTAMP
kubernetes-minion-template  n1-standard-2          2016-12-10T04:25:37.527-
08:00
Created [https://www.googleapis.com/compute/v1/projects/dynamic-nomad-152102/zones/us-central1-b/instanceGroupManagers/kubernetes-minion-group].
NAME                LOCATION      SCOPE  BASE_INSTANCE_NAME  SIZE  TARGET_SIZE  INSTANCE_TEMPLATE  AUTOSCALED
kubernetes-minion-group  us-central1-b  zone  kubernetes-minion-group  0    3
kubernetes-minion-template  no
Waiting for group to become stable, current operations: creating: 3
Waiting for group to become stable, current operations: creating: 3
Waiting for group to become stable, current operations: creating: 1
Group is stable

```

Google Cloud Platform GSW K8s 3

Compute Engine

Instance groups

VM instances

Instance groups

Instance templates

Disks

kubernetes-minion-group

Members Details

Zone: us-central1-b Template: kubernetes-minion-template Autoscaling: Off

CPU utilization

Google Cloud Platform GSW K8s 3

Compute Engine

VM instances

CREATE INSTANCE

IMPORT VM

REFRESH

Filter VM instances

Columns

<input type="checkbox"/>	Name	Zone	Recommendation	Internal IP	External
<input checked="" type="checkbox"/>	kubernetes-master	us-central1-b		10.128.0.2	104.155

Select an instance

LABELS MONITORING

No instances selected.



Your connection is not private

Attackers might be trying to steal your information from **35.184.174.88** (for example, passwords, messages, or credit cards). [Learn more](#)

NET::ERR_CERT_AUTHORITY_INVALID

Automatically send some [system information and page content](#) to Google to help detect dangerous apps and sites. [Privacy policy](#)

HIDE ADVANCED

[Back to safety](#)

This server could not prove that it is **35.184.174.88**; its security certificate is not trusted by your computer's operating system. This may be caused by a misconfiguration or an attacker intercepting your connection.

[Proceed to 35.184.174.88 \(unsafe\)](#)

Kubernetes Dashboard

Kubeconfig

Please select the kubeconfig file that you have created to configure access to the cluster. To find out more about how to configure and use kubeconfig file, please refer to the [Configure Access to Multiple Clusters](#) section.

Token

Every Service Account has a Secret with valid Bearer Token that can be used to log in to Dashboard. To find out more about how to configure and use Bearer Tokens, please refer to the [Authentication](#) section.

Choose kubeconfig file ...

SIGN IN

SKIP

kubernetes [+ CREATE](#)

Overview

Cluster Discovery and Load Balancing

Namespaces

Nodes

Persistent Volumes

Roles

Storage Classes

Namespace default

Overview

Workloads

Cron Jobs

Daemon Sets

Deployments

Jobs

Services

Name	Labels	Cluster IP	Internal endpoints	External endpoints	Age
kubernetes	component: apise... provider: kubernet...	10.0.0.1	kubernetes:443 TCP kubernetes:0 TCP	-	25 minutes

Secrets

Name	Type	Age
default-token-dnx5r	kubernetes.io/service-account-token	24 minutes

kubernetes Nodes [+ CREATE](#)

Admin

Namespaces

Nodes

Persistent Volumes

Namespace default

Workloads

Deployments

Replica Sets

Replication Controllers

Daemon Sets

Pet Sets

Jobs

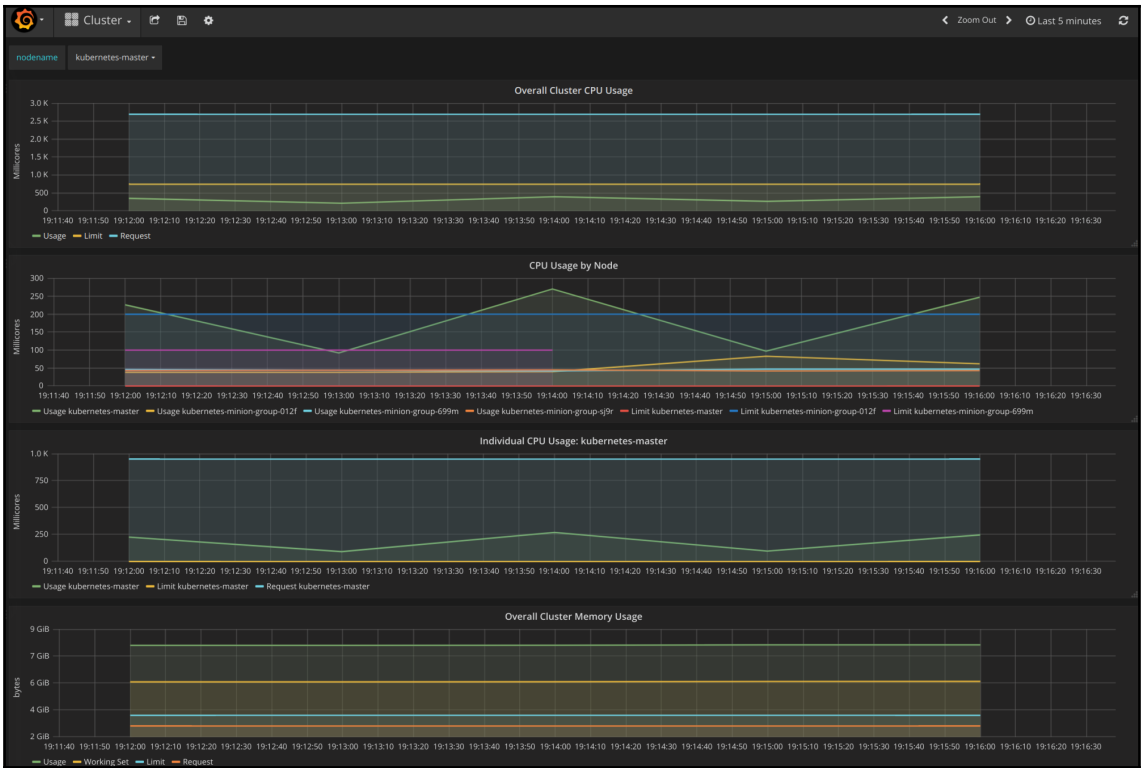
Pods

CPU usage history

Memory usage history

Name	Labels	Ready	Age
kubernetes-master	beta.kubernetes.io/a... beta.kubernetes.io/i... beta.kubernetes.io/o... failure-domain.beta.... failure-domain.beta....	True	9 hours

[show all labels](#)






IMAGE	STATUS
gcr.io/google_containers/node-problem-detector:v0.1	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/fluentd-gcp:1.21	Up 13 hours
gcr.io/google_containers/kube-apiserver:fa481b6112db7dcce46bfc8c8fbf149a2	Up 13 hours
gcr.io/google_containers/etcd:2.2.1	Up 13 hours
gcr.io/google_containers/etcd:2.2.1	Up 13 hours
gcr.io/google_containers/rescheduler:v0.2.1	Up 13 hours
gcr.io/google_containers/glibc:0.8.0	Up 13 hours
gcr.io/google_containers/kube-addon-manager:v5.1	Up 13 hours
gcr.io/google_containers/etcd-empty-dir-cleanup:0.0.1	Up 13 hours
gcr.io/google_containers/kube-controller-manager:9b1fc8f7afac597ccb49e34778214c49	Up 13 hours
gcr.io/google_containers/kube-scheduler:67b73a442b6a6f362a086ea4ab8dc1cd	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours

IMAGE	STATUS
gcr.io/google_containers/exechealthz-amd64:1.2	Up 13 hours
gcr.io/google_containers/kube-dnsmasq-amd64:1.4	Up 13 hours
gcr.io/google_containers/heapster_grafana:v3.1.1	Up 13 hours
gcr.io/google_containers/kubedns-amd64:1.8	Up 13 hours
gcr.io/google_containers/heapster_influxdb:v0.7	Up 13 hours
gcr.io/google_containers/defaultbackend:1.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/fluentd-gcp:1.25	Up 13 hours
gcr.io/google_containers/node-problem-detector:v0.1	Up 13 hours
gcr.io/google_containers/kube-proxy:b87ffd2bf726a72a00bbc021970cb855	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours
gcr.io/google_containers/pause-amd64:3.0	Up 13 hours

Launch Instance ▾
Connect
Actions ▾

<input type="checkbox"/>	Name	Instance ID	Instance Type	Availability Zone	Instance State
<input type="checkbox"/>	master-us-east-2a.masters.gswk8s3.k8s.local	i-04dbc80b39fe7d1da	c4.large	us-east-2a	● running
<input type="checkbox"/>	nodes.gswk8s3.k8s.local	i-04d4f3d43b5165fbc	t2.medium	us-east-2a	● running
<input type="checkbox"/>	nodes.gswk8s3.k8s.local	i-0d1872bc6fe184efd	t2.medium	us-east-2a	● running

```

NAME                                READY   STATUS    RESTARTS   AGE
calico-etcd-7ckip                   1/1    Running   0           43s
calico-node-em9l7                    2/2    Running   0           43s
calico-policy-controller-i43ct      1/1    Running   0           43s
dummy-2088944543-efrgw              1/1    Running   0           2m
etcd-ip-172-30-0-26                  1/1    Running   0           1m
kube-apiserver-ip-172-30-0-26        1/1    Running   0           2m
kube-controller-manager-ip-172-30-0-26 1/1    Running   0           2m
kube-discovery-1150918428-1kntn     1/1    Running   0           2m
kube-dns-654381707-6u52r            2/3    Running   0           1m
kube-proxy-00wu7                     1/1    Running   0           1m
kube-scheduler-ip-172-30-0-26        1/1    Running   0           1m
info: 1 completed object(s) was(were) not shown in pods list. Pass --show-all
to see all objects.

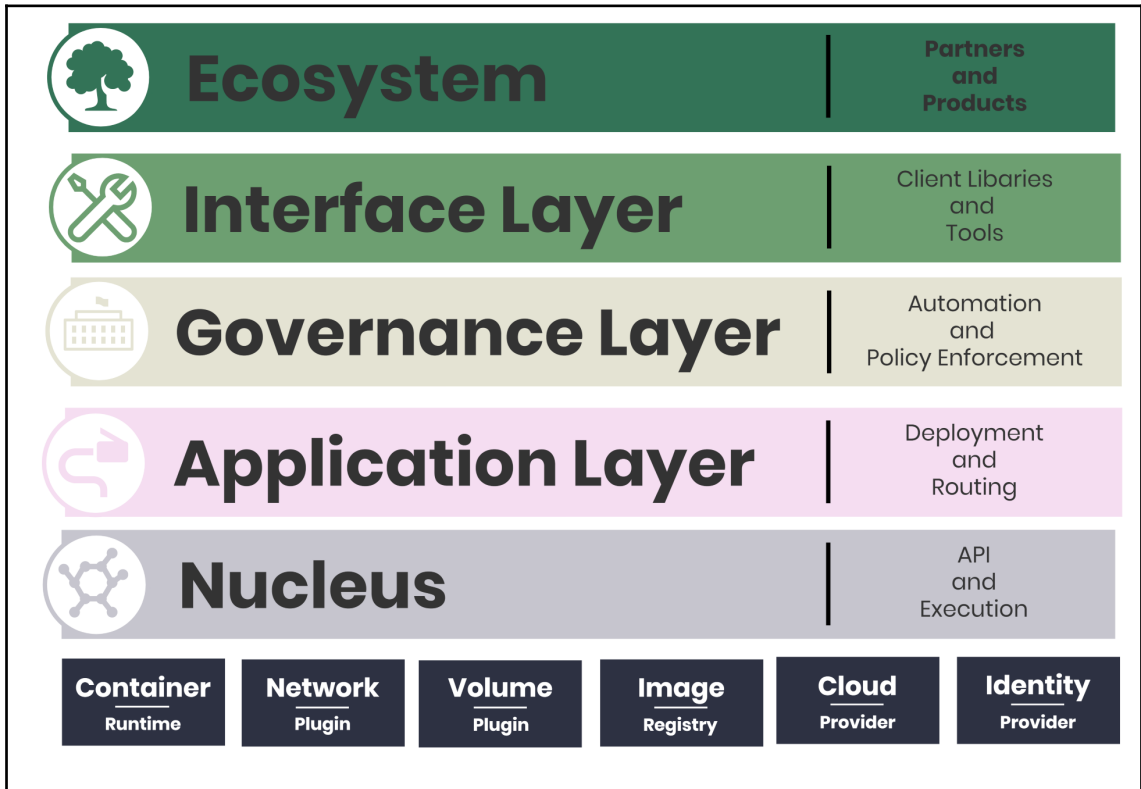
```

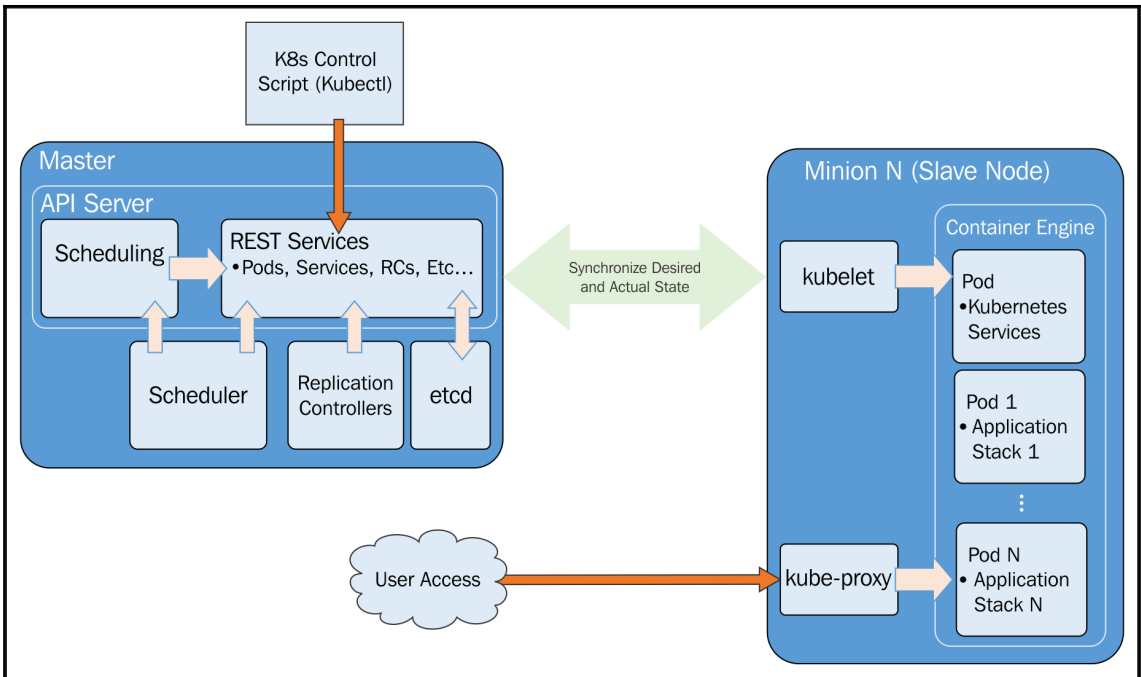
```

NAME                                STATUS    AGE
ip-172-30-0-22                      Ready    6m
ip-172-30-0-26                      Ready, master 8m
ip-172-30-0-28                      Ready    6m
ip-172-30-0-8                       Ready    6m

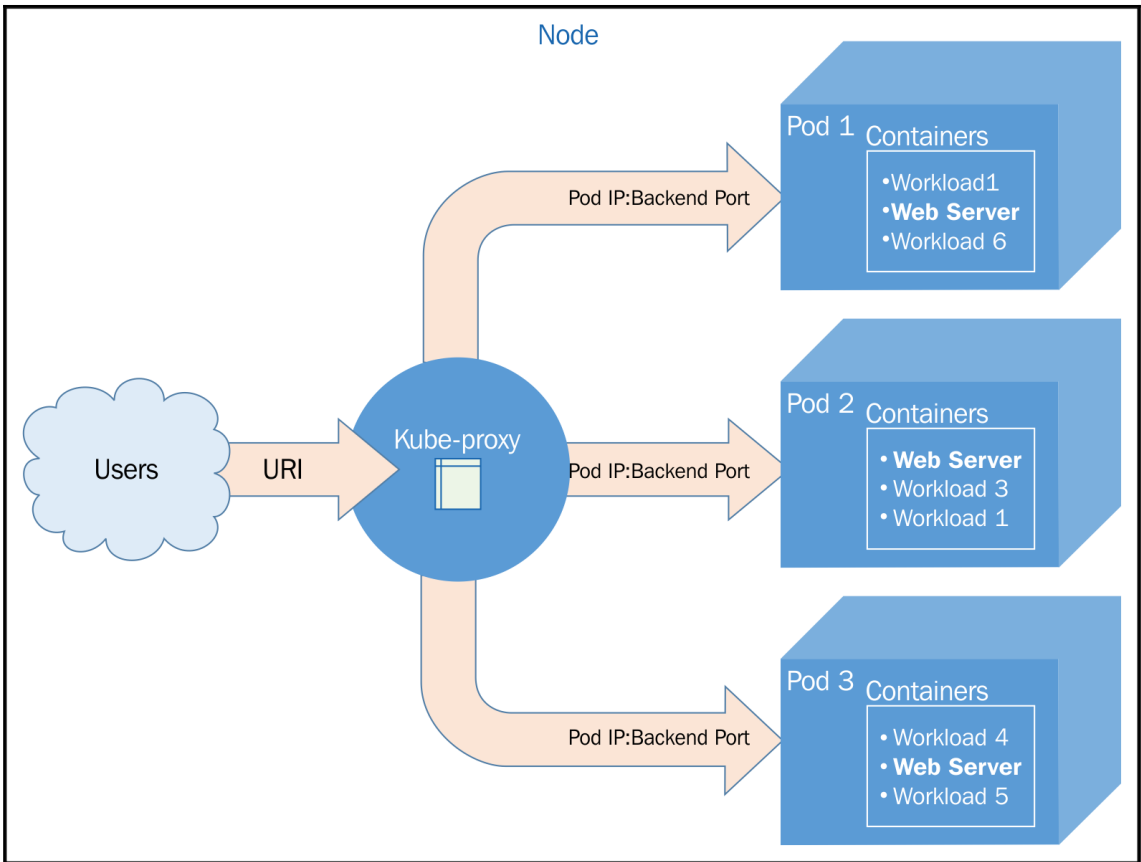
```

Chapter 2: Building a Foundation with Core Kubernetes Constructs





Node Condition	Description
OutOfDisk	True if there is insufficient free space on the node for adding new pods, otherwise False
Ready	True if the node is healthy and ready to accept pods, False if the node is not healthy and is not accepting pods, and Unknown if the node controller has not heard from the node in the last 40 seconds
MemoryPressure	True if pressure exists on the node memory – that is, if the node memory is low; otherwise False
DiskPressure	True if pressure exists on the disk size – that is, if the disk capacity is low; otherwise False
NetworkUnavailable	True if the network for the node is not correctly configured, otherwise False
ConfigOK	True if the kubelet is correctly configured, otherwise False



NAME	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
kubernetes	10.0.0.1	<none>	443/TCP	11m
node-js	10.0.200.192	35.184.181.18	80:30874/TCP	4m

```

Host: node-js-u26fd
Running OS: linux
Uptime: 525274
Network Information: 10.244.1.17, fe80::42:aff:fef4:111
DNS Servers: 10.0.0.10,169.254.169.254,10.240.0.1

```

NAME	READY	STATUS	RESTARTS	AGE
node-js-1fxoy	1/1	Running	0	1d
node-js-m4w4a	1/1	Running	0	1d
node-js-sjc03	1/1	Running	0	1d

```

Name:                node-js-sjc03
Namespace:          default
Image(s):           petegoo/node-express-sample:latest
Node:               kubernetes-minion-aqdf/10.240.142.178
Labels:             name=node-js
Status:             Running
Reason:
Message:
IP:                 10.244.0.10
Replication Controllers:  node-js (3/3 replicas created)
Containers:
  node-js:
    Image:           petegoo/node-express-sample:latest
    Limits:
      cpu:           100m
    State:           Running
      Started:       Tue, 28 Jul 2015 16:57:33 -0400
      Ready:         True
      Restart Count: 0
Conditions:
  Type      Status
  Ready     True
No events.

```


☰ **kubernetes** Nodes > gke-cluster-1-default-pool-3185750f-q6sx + CREATE

Admin

Namespaces

Nodes

Persistent Volumes

Namespace

default ▾

Workloads

Deployments

Replica Sets

Replication Controllers

Daemon Sets

Pet Sets

Jobs

Pods

Services and discovery

Services

Events

Message	Source	Sub-object	Count	First seen	Last seen
Starting kubelet.	kubelet gke-cluster-1-default-pool-3185750f-q6sx	-	1	22/12/16 21:42 UTC	22/12/16 21:42 UTC
Node gke-cluster-1-default-pool-3185750f-q6sx status is now: NodeHasSufficientDisk	kubelet gke-cluster-1-default-pool-3185750f-q6sx	-	17	22/12/16 21:42 UTC	22/12/16 21:44 UTC
Node gke-cluster-1-default-pool-3185750f-q6sx status is now: NodeHasSufficientMemory	kubelet gke-cluster-1-default-pool-3185750f-q6sx	-	17	22/12/16 21:42 UTC	22/12/16 21:44 UTC
Node gke-cluster-1-default-pool-3185750f-q6sx status is now: NodeHasNoDiskPressure	kubelet gke-cluster-1-default-pool-3185750f-q6sx	-	17	22/12/16 21:42 UTC	22/12/16 21:44 UTC

NAME	DESIRED	CURRENT	READY	AGE
node-js-labels	3	3	3	46s

NAME	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
kubernetes	10.0.0.1	<none>	443/TCP	5d

NAME	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
node-js	10.0.13.62	104.197.124.230	80:30798/TCP	14h
node-js-labels	10.0.207.25	104.154.54.104	80:31315/TCP	1m

```
Name:          node-js
Namespace:    default
Image(s):     jonbaier/node-express-info:latest
Selector:     name=node-js
Labels:       name=node-js
Replicas:     3 current / 3 desired
Pods Status:  3 Running / 0 Waiting / 0 Succeeded / 0 Failed
No volumes.
Events:
  FirstSeen    LastSeen      Count   From              SubobjectPath    Type
  Reason      Message
  -----
--
  42s         42s          1      {replication-controller }
SuccessfulCreate   Created pod: node-js-7esbp
  42s         42s          1      {replication-controller }
SuccessfulCreate   Created pod: node-js-istu0
  42s         42s          1      {replication-controller }
SuccessfulCreate   Created pod: node-js-im7jw
```

```


Name:          node-js-7esbp
Namespace:    default
Node:         kubernetes-minion-group-k0rn/10.128.0.3
Start Time:   Mon, 02 Jan 2017 13:54:22 -0500
Labels:       name=node-js
Status:       Running
IP:           10.244.1.18
Controllers:  ReplicationController/node-js
Containers:
  node-js:
    Container ID:  docker://ce35e1fba7c3464cc89607ebd335250a7b52bebd5e03683e3f6313f35fe68244
    Image:         jonbaier/node-express-info:latest
    Image ID:     docker://sha256:6a276384568844d1840049552f79c69311c3132d3a2b884a3e9c4e51087a436b
    Port:         80/TCP
    Requests:
      CPU:         100m
    State:        Waiting
      Reason:      CrashLoopBackOff
    Last State:   Terminated
      Reason:      Error
      Exit Code:   137
    Started:      Mon, 02 Jan 2017 14:13:42 -0500
    Finished:     Mon, 02 Jan 2017 14:14:42 -0500
    Ready:        False
    Restart Count: 9
    Liveness:     http-get http://:80/status/ delay=30s timeout=1s period=10s #success=1 #failure=3
  volume mounts:
    /var/run/secrets/kubernetes.io/serviceaccount from default-token-7z353 (ro)
Environment Variables: <none>
Conditions:
  Type           Status
  Initialized    True
  Ready          False
  PodScheduled   True
Volumes:
  default-token-7z353:
    Type:          Secret (a volume populated by a Secret)
    SecretName:    default-token-7z353
QoS Class:       Burstable
Tolerations:     <none>
Events:
  FirstSeen      LastSeen        Count   From              SubobjectPath
  -----
  22m            22m             1      {default-scheduler}  Normal
  Successfully assianed node-js-7esbp to kubernetes-minion-group-k0rn
  21m            21m             1      {kubelet kubernetes-minion-group-k0rn} spec.containers{node-js}
  Created        Created container with docker id 4b2b5587a119; Security:[seccomp=unconfined]
  21m            21m             1      {kubelet kubernetes-minion-group-k0rn} spec.containers{node-js}
  Started        Started container with docker id 4b2b5587a119
  20m            20m             1      {kubelet kubernetes-minion-group-k0rn} spec.containers{node-js}
  Killing        Killing container with docker id 4b2b5587a119; pod "node-js-7esbp_default(df9e1d36-d11c-11e6-9141-42010a800002)" container "node-js" is unhealthy, it will be killed and re-created.
  20m            20m             1      {kubelet kubernetes-minion-group-k0rn} spec.containers{node-js}
  Created        Created container with docker id 53e4c1ec9e20; Security:[seccomp=unconfined]
  20m            20m             1      {kubelet kubernetes-minion-group-k0rn} spec.containers{node-js}
  Started        Started container with docker id 53e4c1ec9e20

```


Nodes

+ CREATE

CPU usage history



Memory usage history



Name	Labels	Ready	Age
✓ kubernetes-master	<ul style="list-style-type: none">beta.kubernetes.io/arch:...beta.kubernetes.io/insta...beta.kubernetes.io/os: li...failure-domain.beta.kub...failure-domain.beta.kub... show all labels	True	6 days
✓ kubernetes-minion-group-7...	<ul style="list-style-type: none">beta.kubernetes.io/arch:...beta.kubernetes.io/insta...beta.kubernetes.io/os: li...failure-domain.beta.kub...failure-domain.beta.kub... show all labels	True	3 days
✓ kubernetes-minion-group-9...	<ul style="list-style-type: none">beta.kubernetes.io/arch:...beta.kubernetes.io/insta...beta.kubernetes.io/os: li...failure-domain.beta.kub...	True	3 days

```

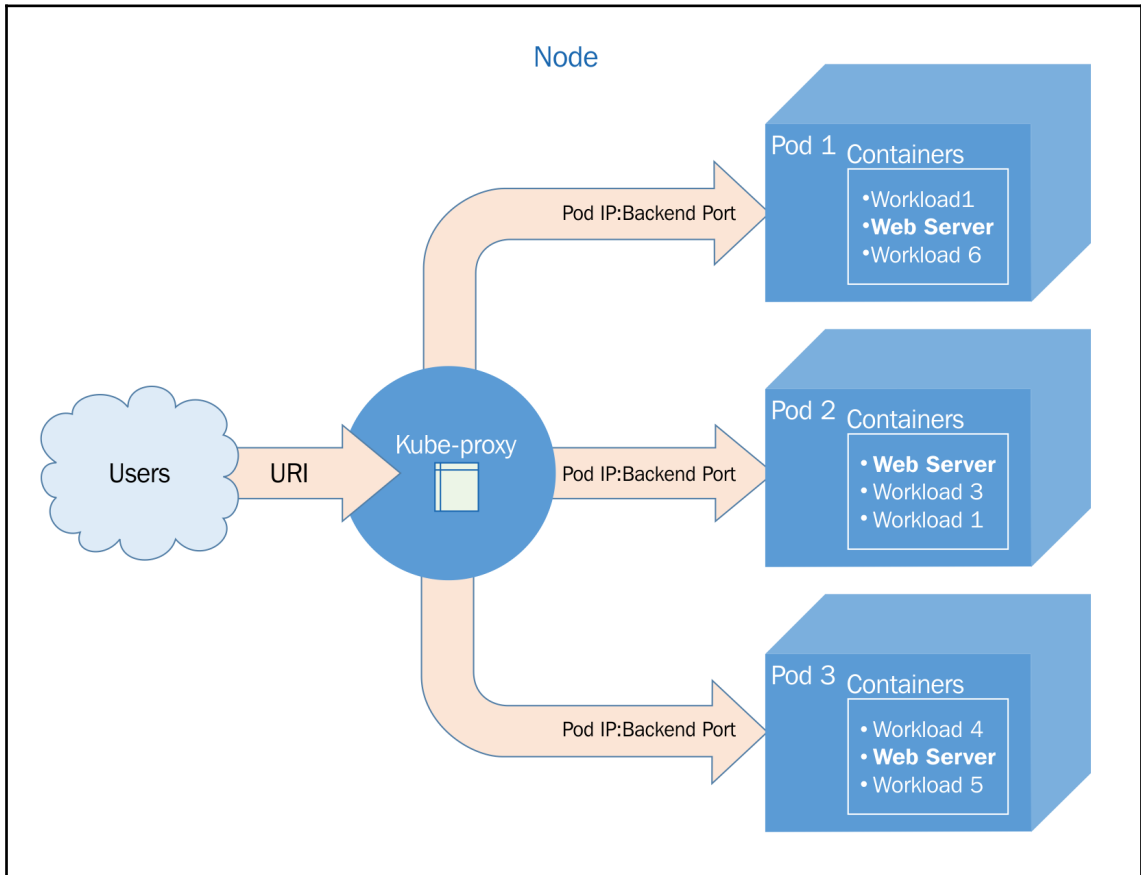
Name:          node-js-constraints-n9dlx
Namespace:    default
Node:         /
Labels:       name=node-js-constraints
Status:       Pending
IP:
Controllers:  ReplicationController/node-js-constraints
Containers:
  node-js-constraints:
    Image:      jonbaier/node-express-info:latest
    Port:       80/TCP
    Limits:
      cpu:       1500m
      memory:    512Mi
    Requests:
      cpu:       1500m
      memory:    512Mi
    Volume Mounts:
      /var/run/secrets/kubernetes.io/serviceaccount from default-token-7z353 (ro)
    Environment Variables:
      <none>
Conditions:
  Type          Status
  PodScheduled  False
Volumes:
  default-token-7z353:
    Type:        Secret (a volume populated by a Secret)
    SecretName:  default-token-7z353
QoS Class:      Guaranteed
Tolerations:    <none>
Events:
  FirstSeen    LastSeen      Count    From          SubobjectPath  Type  Reason
  -----
  1m           1m            3        {default-scheduler }
FailedScheduling pod (node-js-constraints-n9dlx) failed to fit in any node
fit failure on node (kubernetes-minion-group-9zf7): Insufficient cpu
fit failure on node (kubernetes-minion-group-k0rn): Insufficient cpu
fit failure on node (kubernetes-minion-group-7th4): Insufficient cpu

  1m           1m            3        {default-scheduler }          Warning FailedScheduling
pod (node-js-constraints-n9dlx) failed to fit in any node
fit failure on node (kubernetes-minion-group-7th4): Insufficient cpu
fit failure on node (kubernetes-minion-group-9zf7): Insufficient cpu
fit failure on node (kubernetes-minion-group-k0rn): Insufficient cpu

  1m           41s           2        {default-scheduler }          Warning FailedScheduling
pod (node-js-constraints-n9dlx) failed to fit in any node
fit failure on node (kubernetes-minion-group-k0rn): Insufficient cpu
fit failure on node (kubernetes-minion-group-7th4): Insufficient cpu
fit failure on node (kubernetes-minion-group-9zf7): Insufficient cpu

```

Chapter 3: Working with Networking, Load Balancers, and Ingress




```
Name:          node-js-labels
Namespace:     default
Labels:        app=node-js-express,deployment=test,name=node-js-labels
Selector:      app=node-js-express,name=node-js-labels
Type:          LoadBalancer
IP:            10.0.115.200
LoadBalancer Ingress: 146.148.56.25
Port:          <unnamed>      80/TCP
NodePort:      <unnamed>      30237/TCP
Endpoints:     10.244.0.29:80,10.244.2.34:80,10.244.2.35:80
Session Affinity: None
No events.
```

NAME	LABELS	SELECTOR	IP(S)	PORT(S)
node-js-internal	name=node-js-internal	name=node-js	10.0.5.134	80/TCP

You have exposed your service on an external port on all nodes in your cluster. If you want to expose this service to the external internet, you may need to set up firewall rules for the service port(s) (tcp:30001) to serve traffic.

See <http://releases.k8s.io/HEAD/docs/user-guide/services-firewalls.md> for more details.
services/node-js-nodeport



Create a new firewall rule

Name [?]

Description (Optional)

Network [?]

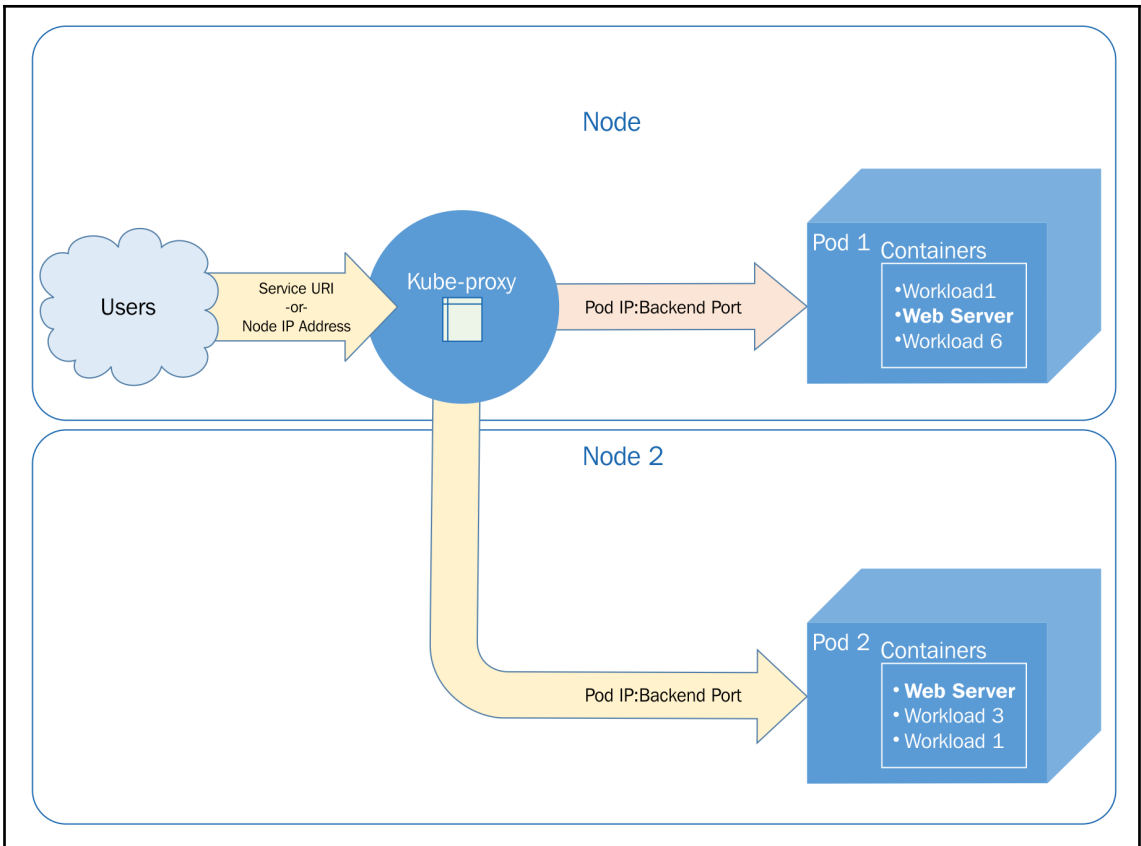
Source filter [?]

Source IP Ranges [?]

Allowed protocols and ports [?]

Target tags (Optional) [?]

Equivalent [REST](#) or [command line](#)



NAME	DESIRED	CURRENT	READY	AGE
kube-dns-v20	1	1	1	8d
kubernetes-dashboard-v1.4.0	1	1	1	8d
l7-default-backend-v1.0	1	1	1	8d
monitoring-influxdb-grafana-v4	1	1	1	8d

NAME	HOSTS	ADDRESS	PORTS	AGE
whale-ingress	a.whale.hey,b.whale.hey	130.211.24.177	80	3h


```

Name:    test
Labels:  <none>
Status:  Active

Resource Quotas
Resource          Used    Hard
---             ---    ---
pods              0      3
replicationcontrollers 0      1
services          0      1

No resource limits.

```

```

Name:          busybox-ns
Namespace:     test
Image(s):      busybox
Selector:      name=busybox-ns
Labels:        name=busybox-ns
Replicas:      3 current / 4 desired
Pods Status:   3 Running / 0 Waiting / 0 Succeeded / 0 Failed
Events:
  FirstSeen          LastSeen          Count   F
rom                SubobjectPath    Reason      Message
  Mon, 17 Aug 2015 16:29:43 -0400    Mon, 17 Aug 2015 16:29:43 -0400    1      {
replication-controller }      successfulCreate      Created p
od: busybox-ns-spfrn
  Mon, 17 Aug 2015 16:29:43 -0400    Mon, 17 Aug 2015 16:29:43 -0400    1      {
replication-controller }      successfulCreate      Created p
od: busybox-ns-xjf6q
  Mon, 17 Aug 2015 16:29:43 -0400    Mon, 17 Aug 2015 16:29:43 -0400    1      {
replication-controller }      successfulCreate      Created p
od: busybox-ns-zeuuy
  Mon, 17 Aug 2015 16:29:44 -0400    Mon, 17 Aug 2015 16:33:01 -0400    18     {
replication-controller }      failedCreate          Error cre
ating: Pod "busybox-ns-" is forbidden: Limited to 3 pods

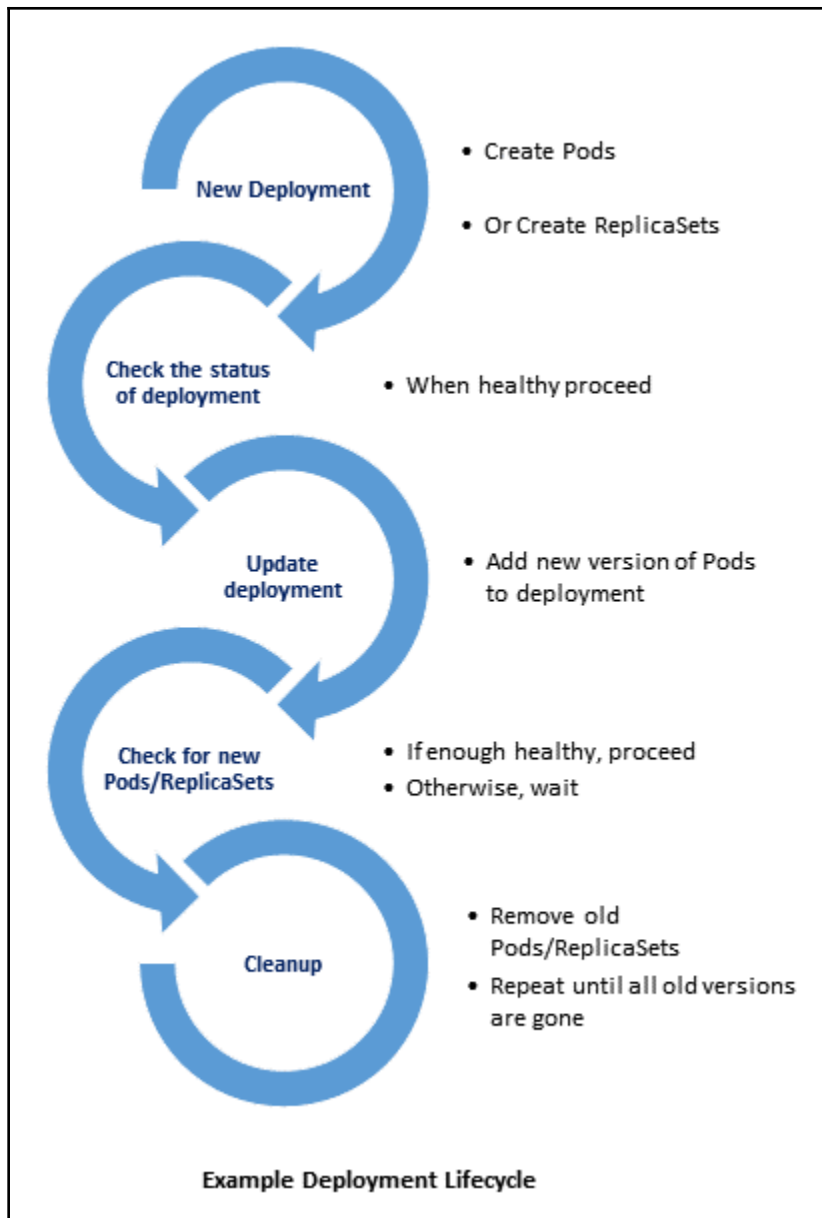
```

Chapter 4: Implementing Reliable Container-Native Applications

NAME	READY	STATUS	RESTARTS	AGE
node-js-deploy-1713031517-itnwi	1/1	Running	0	6m
node-js-deploy-1713031517-nx8vs	1/1	Running	0	6m
node-js-deploy-1713031517-uge5y	1/1	Running	0	6m

Image: jonbaier/pod-scaling:0.1

NAME	READY	STATUS	RESTARTS	AGE
node-js-deploy-1794296158-5wivi	1/1	Running	0	5m
node-js-deploy-1794296158-b2any	1/1	Running	0	5m
node-js-deploy-1794296158-y2tx3	1/1	Running	0	5m



```

REVISION      CHANGE-CAUSE
1             kubectl scale deployment node-js-deploy --replicas 3
2             kubectl set image deployment/node-js-deploy node-js-
deploy=jonbaier/pod-scaling:0.2
3             kubectl set image deployment/node-js-deploy node-js-
deploy=jonbaier/pod-scaling:0.3

```

NAME	READY	STATUS	RESTARTS	AGE
node-js-deploy-1875560799-ehi0o	1/1	Running	0	40m
node-js-deploy-1875560799-tqset	1/1	Running	0	40m
node-js-deploy-1907673490-cadw2	0/1	ErrImagePull	0	1m
node-js-deploy-1907673490-qvc9w	0/1	ErrImagePull	0	1m

```

REVISION      CHANGE-CAUSE
1             kubectl scale deployment node-js-deploy --replicas 3
2             kubectl set image deployment/node-js-deploy node-js-de
ploy=jonbaier/pod-scaling:0.2
4             kubectl set image deployment/node-js-deploy node-js-de
ploy=jonbaier/pod-scaling:42.0
5             kubectl set image deployment/node-js-deploy node-js-de
ploy=jonbaier/pod-scaling:0.3

```

NAME	REFERENCE	TARGET	CURRENT	MIN
PODS MAXPODS	AGE			
node-js-deploy 6	Deployment/node-js-deploy 3h	10%	0%	3
node-js-scale 3	ReplicationController/node-js-scale 10d	30%	0%	1

```

$ kubectl get hpa
NAME           REFERENCE                TARGET    CURRENT    MIN
PODS  MAXPODS  AGE
node-js-deploy  Deployment/node-js-deploy  10%      20%      3
6
node-js-scale   ReplicationController/node-js-scale  30%      0%      1
3
9d

$ kubectl get deploy
NAME           DESIRED    CURRENT    UP-TO-DATE    AVAILABLE    AGE
boomload-deploy  1          1          1              1            4m
node-js-deploy  6          6          6              6            10d

```

```

Name: long-task
Namespace: default
Image(s): docker/whalesay
Selector: controller-uid=eff2fcd2-d5e1-11e6-90ee-42010a800002
Parallelism: 1
Completions: 1
Start Time: Sun, 08 Jan 2017 15:35:05 -0500
Labels: <none>
Pods Statuses: 0 Running / 1 Succeeded / 0 Failed
No volumes.
Events:
  FirstSeen      LastSeen        Count   From              SubobjectPath  Type
  Reason          Message
  -----
  4m              4m              1      {job-controller }  SuccessfulCreate  Normal
  Created pod: long-task-a6i9v

```

Logs from long-task in long-task-a6i9v A Tt

```

2017-01-08T20:35:19.747705449Z -----
2017-01-08T20:35:19.747740036Z < Finishing that task in a jiffy >
2017-01-08T20:35:19.747747931Z -----
2017-01-08T20:35:19.747752786Z \
2017-01-08T20:35:19.747756904Z \
2017-01-08T20:35:19.747761135Z \
2017-01-08T20:35:19.747765442Z ##
2017-01-08T20:35:19.747770620Z ## ## ## ==
2017-01-08T20:35:19.747775118Z ## ## ## ## ===
2017-01-08T20:35:19.747779445Z /"-----"___/ ===
2017-01-08T20:35:19.747784865Z ~~~ {~ ~~~~ ~~~ ~~~~ ~ ~ / ===- ~~~
2017-01-08T20:35:19.747789250Z \----- o ---/
2017-01-08T20:35:19.747793504Z \
2017-01-08T20:35:19.747797946Z \-----\

```

Logs from 1/8/17 3:35 PM to 1/8/17 3:35 PM |< < > >|

NAME	DESIRED	CURRENT	NODE-SELECTOR	AGE
node-problem-detector-v0.1	4	4	<none>	13d

```
Name:          kubernetes-minion-group-1l6g
Labels:        beta.kubernetes.io/arch=amd64
               beta.kubernetes.io/instance-type=n1-standard-2
               beta.kubernetes.io/os=linux
               failure-domain.beta.kubernetes.io/region=us-central1
               failure-domain.beta.kubernetes.io/zone=us-central1-b
               kubernetes.io/hostname=kubernetes-minion-group-1l6g
Taints:        <none>
CreationTimestamp:  Wed, 11 Jan 2017 07:48:16 -0500
Phase:
Conditions:
```

Chapter 5: Exploring Kubernetes Storage Concepts

```
/home/k8s/nodejs# kubectl.sh exec memory-pd -- ls -lh | grep memory  
drwxrwxrwt 2 root root 40 Oct 24 15:21 memory-pd
```

- Home
- Permissions
- APIs & auth
- Monitoring
 - Traces
 - Logs
 - Dashboards & alerts
- Source Code
- Cloud Launcher
- Deployments
- Compute
 - App Engine
 - Compute Engine
 - VM instances
 - Instance groups
 - Instance templates
 - Disks
 - Snapshots
 - Images
 - Metadata
 - Health checks
 - Zones
 - Operations
 - Quotas
 - Settings
 - Container Engine
- Networking
- Storage
 - Cloud Bigtable
 - Cloud Datastore
 - Cloud SQL

←

Create a new disk

Name ?

Description (Optional)

Zone ?

Disk Type ?

Source type ?

Size (GB) ?

i You have entered a volume size of under 200 GB. This may result in reduced performance. [Learn more](#)

Estimated performance ?

Operation Type	Read	Write
Sustained random IOPS limit	3	15
Sustained throughput limit (MB/s)	1.2	0.9

Encryption ?

Equivalent [REST](#) or [command line](#)

```
Name:          test-gce
Namespace:     default
Node:          kubernetes-minion-group-zwpm/10.128.0.4
Start Time:   Sun, 15 Jan 2017 16:51:02 -0500
Labels:       <none>
Status:       Running
IP:           10.244.4.5
Controllers:  <none>
Containers:
  test-gce:
    Container ID:  docker://15871d81eb72557cc230df70a5c724617289d710a550da66e4dfaf7083
    Image:         nginx:latest
    Image ID:     docker://sha256:01f818af747d88b4ebca7cdabd0c581e406e0e790be72678d25
    Port:         80/TCP
    Requests:
      cpu:        100m
    State:        Running
      Started:    Sun, 15 Jan 2017 16:53:00 -0500
    Ready:        True
    Restart Count: 0
    Volume Mounts:
      /usr/share/nginx/html from gce-pd (rw)
      /var/run/secrets/kubernetes.io/serviceaccount from default-token-728d1 (ro)
    Environment Variables: <none>
Conditions:
  Type           Status
  Initialized    True
  Ready          True
  PodScheduled  True
Volumes:
  gce-pd:
    Type:          GCEPersistentDisk (a Persistent Disk resource in Google Compute Engine)
    PDName:       mysite-volume-1
    FSType:       ext4
    Partition:    0
    ReadOnly:     false
  default-token-728d1:
    Type:          Secret (a volume populated by a Secret)
    SecretName:   default-token-728d1
QoS Class:      Burstable
Tolerations:    <none>
```

```

Name:          http-pd
Namespace:    default
Labels:       name=http-pd
Selector:     name=http-pd
Type:         LoadBalancer
IP:           10.0.118.195
LoadBalancer Ingress: 130.211.186.84
Port:         http      80/TCP
NodePort:     http      32429/TCP
Endpoints:    10.244.2.15:80,10.244.2.16:80,10.244.3.5:80
Session Affinity: None
No events.

```

NAME	DESIRED	CURRENT	AGE
whaleset	3	3	46s

NAME	READY	STATUS	RESTARTS	AGE
whaleset-0	1/1	Running	0	54s
whaleset-1	1/1	Running	0	29s
whaleset-2	0/1	ContainerCreating	0	11s

NAME	STATUS	VOLUME	CAPACITY	ACCESSMODES
AGE				
www-whaleset-0 4m	Bound	pvc-43346a3d-e024-11e6-af6d-42010a800002	1Gi	RWO
www-whaleset-1 4m	Bound	pvc-43381dc9-e024-11e6-af6d-42010a800002	1Gi	RWO
www-whaleset-2 4m	Bound	pvc-433a3864-e024-11e6-af6d-42010a800002	1Gi	RWO

NAME	REASON	AGE	CAPACITY	ACCESSMODES	RECLAIMPOLICY	STATUS
CLAIM						
pvc-43346a3d-e024-11e6-af6d-42010a800002		4m	1Gi	RWO	Delete	Bound
default/www-whaleset-0		4m				
pvc-43381dc9-e024-11e6-af6d-42010a800002		4m	1Gi	RWO	Delete	Bound
default/www-whaleset-1		4m				
pvc-433a3864-e024-11e6-af6d-42010a800002		4m	1Gi	RWO	Delete	Bound
default/www-whaleset-2		4m				

Chapter 6: Application Updates, Gradual Rollouts, and Autoscaling

```
Creating node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6
At beginning of loop: node-js-scale replicas: 2, node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 replicas: 1
Updating node-js-scale replicas: 2, node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 replicas: 1
At end of loop: node-js-scale replicas: 2, node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 replicas: 1
At beginning of loop: node-js-scale replicas: 1, node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 replicas: 2
Updating node-js-scale replicas: 1, node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 replicas: 2
At end of loop: node-js-scale replicas: 1, node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 replicas: 2
At beginning of loop: node-js-scale replicas: 0, node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 replicas: 3
Updating node-js-scale replicas: 0, node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 replicas: 3
At end of loop: node-js-scale replicas: 0, node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 replicas: 3
Update succeeded. Deleting old controller: node-js-scale
Renaming node-js-scale-10ea08ff9a118ac6a93f85547ed2d8f6 to node-js-scale
node-js-scale
```


Pod Scaling v0.1

Host: node-js-scale-1w562

Running OS: linux

Uptime: 1332

Network Information: 10.244.1.7, fe80::989b:9cff:fe60:d933

DNS Servers: 10.0.0.10

Pod Scaling v0.2

Host: node-js-scale-3ad907156054d5840a726bcb4edb7cbf-pxr0v

Running OS: linux

Uptime: 3071

Network Information: 10.244.0.6, fe80::4829:99ff:feff:79b7

DNS Servers: 10.0.0.10

NAME	REFERENCE	TARGET	CURRENT
MINPODS	MAXPODS		
node-js-scale	ReplicationController/node-js-scale	30%	0%
1	3		
	AGE		
	2d		

NAME	REFERENCE	TARGET	CURRENT
MINPODS	MAXPODS		
node-js-scale	ReplicationController/node-js-scale	30%	49%
1	3		
	AGE		
	2d		

- Allow HTTP traffic
- Allow HTTPS traffic

Availability policies

Preemptibility	Off (recommended)
Automatic restart	On (recommended)
On host maintenance	Migrate VM instance (recommended)

Custom metadata

```
kube-  
env  ENV_TIMESTAMP: '2015-07-26T13:42:41+0000'  
      INSTANCE_PREFIX: 'kubernetes'  
      NODE_INSTANCE_PREFIX: 'kubernetes-minion'  
      CLUSTER_IP_RANGE: '10.244.0.0/16'  
      SERVER_BINARY_TAR_URL: 'https://storage.googleapis.com/kubernetes-staging-f8a93094f0/devel/kubernetes-server-linux-amd64.tar.gz'  
      SERVER_BINARY_TAR_HASH: 'b2968ede4437bbc6aeb2ca84cf26e01fb20ec988'  
      SALT_TAR_URL: 'https://storage.googleapis.com/kubernetes-staging-f8a93094f0/devel/kubernetes-salt.tar.gz'  
      SALT_TAR_HASH: '434740483205e0a755f6806574787e3d639123f4'  
      SERVICE_CLUSTER_IP_RANGE: '10.0.0.0/16'  
      KUBERNETES_MASTER_NAME: 'kubernetes-master'  
      ALLOCATE_NODE_CIDRS: 'true'  
      ENABLE_CLUSTER_MONITORING: 'googleinfluxdb'  
      ENABLE_CLUSTER_LOGGING: 'true'  
      ENABLE_NODE_LOGGING: 'true'  
      LOGGING_DESTINATION: 'gcp'  
      ELASTICSEARCH_LOGGING_REPLICAS: '1'  
      ENABLE_CLUSTER_DNS: 'true'  
      DNS_REPLICAS: '1'  
      DNS_SERVER_IP: '10.0.0.10'  
      DNS_DOMAIN: 'cluster.local'  
      KUBELET_TOKEN: 'E6OZsbuQr0efOGJDMIsY59xY4DyjkjXK'  
      KUBE_PROXY_TOKEN: '1cOJI6Tb2hOjBxis0bI8xwI6OaktUd9A'
```

← Instance groups ✎ EDIT GROUP 🗑 DELETE GROUP

🟢 kubernetes-minion-group

Members Details

Zone: **us-central1-b** Template: **kubernetes-minion-template** Autoscaling: **Off** In use by:

CPU utilization ▾ 1 hour 6h 12h 1 day 2d 4d 7d 14d 30d

CPU

% CPU

Jan 23, 6:00 PM Jan 23, 6:15 PM Jan 23, 6:30 PM Jan 23, 6:47 PM

■ CPU: 3.209

<input type="checkbox"/> Name ^	Disk	External IP	Connect
<input type="checkbox"/> ✔ kubernetes-minion-group-6q7r	kubernetes-minion-group-6q7r, kubernetes-dynamic-pvc-f20c86a0-e0c1-11e6-8dba-42010a800002	104.197.92.178	SSH ▾
<input type="checkbox"/> ✔ kubernetes-minion-group-83bw	kubernetes-minion-group-83bw, kubernetes-dynamic-pvc-f205ca42-e0c1-11e6-8dba-42010a800002	104.198.157.91	SSH ▾
<input type="checkbox"/> ✔ kubernetes-minion-group-xgtg	kubernetes-minion-group-xgtg, kubernetes-dynamic-pvc-f20a0238-e0c1-11e6-8dba-42010a800002	130.211.170.182	SSH ▾

← Instance groups EDIT GROUP DELETE GROUP

Edit kubernetes-minion-group

Zone
us-central1-b

[Specify port name mapping](#) (Optional)

Instance template ?
kubernetes-minion-template

Autoscaling ?
Off

Number of instances
3

Autohealing
VMs in the group are recreated as needed. You can use a health check to recreate a VM if the health check finds the VM unresponsive. If you do not select a health check, VMs are recreated only when stopped. [Learn more](#)

Health check
No health check

Initial delay ?
300 seconds

Save Cancel

[Create Auto Scaling group](#)
Actions ▾

Filter:

 << < 1 to 1 of 1 Auto Scaling Groups > >>

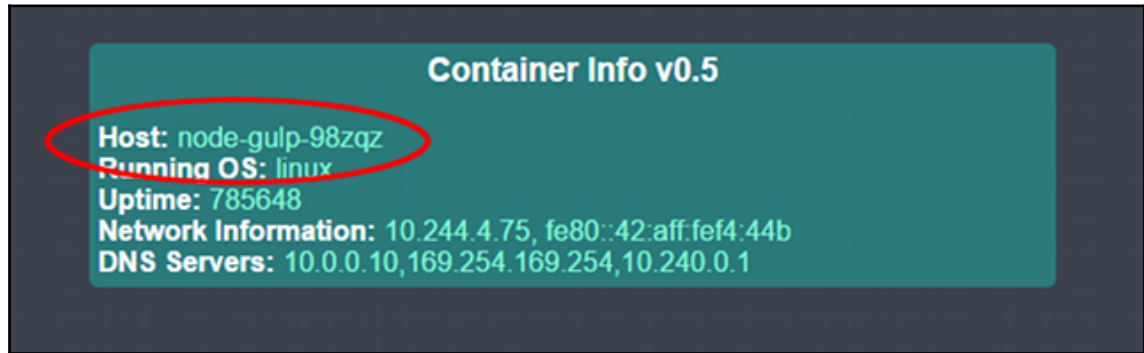
<input type="checkbox"/>	Name	Launch Configuration	Instances	Desired	Min	Max	Availability Zones
<input checked="" type="checkbox"/>	kubernetes-mi...	kubernetes-minion-group	4	4	4	4	us-west-2a

Auto Scaling Group: **kubernetes-minion-group**

[Details](#)
[Activity History](#)
[Scaling Policies](#)
[Instances](#)
[Notifications](#)
[Tags](#)
[Edit](#)

Launch Configuration	kubernetes-minion-group		
Load Balancers			
Desired	4	Availability Zone(s)	us-west-2a
Min	4	Subnet(s)	subnet-c66eb4b1
Max	4	Default Cooldown	300
Health Check Type	EC2	Placement Group	
Health Check Grace Period	0	Suspended Processes	
Termination Policies	Default	Enabled Metrics	
Creation Time	Sun Oct 25 12:08:06 GMT-400 2015		

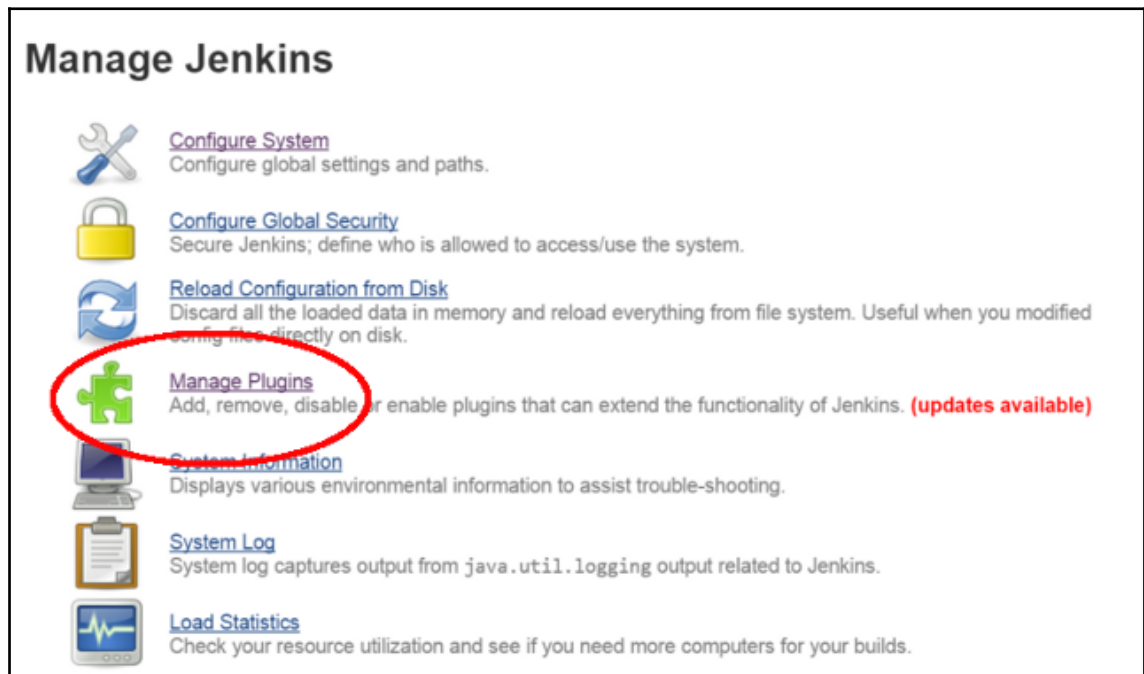
Chapter 7: Designing for Continuous Integration and Delivery



A screenshot of the Jenkins 'Container Info v0.5' page. The background is dark grey, and the text is white. A red oval highlights the first three lines of information: 'Host: node-gulp-98zqz', 'Running OS: linux', and 'Uptime: 785648'. Below these are 'Network Information: 10.244.4.75, fe80::42:aff:fe4:44b' and 'DNS Servers: 10.0.0.10, 169.254.169.254, 10.240.0.1'.








Container Info v0.5

Host: node-gulp-98zqz
Running OS: linux
Uptime: 785648
Network Information: 10.244.4.75, fe80::42:aff:fe4:44b
DNS Servers: 10.0.0.10, 169.254.169.254, 10.240.0.1



A screenshot of the Jenkins 'Manage Jenkins' page. The page has a white background with a list of management options, each with an icon and a description. A red oval highlights the 'Manage Plugins' option, which includes the text '(updates available)' in red. The other options are 'Configure System', 'Configure Global Security', 'Reload Configuration from Disk', 'System Information', 'System Log', and 'Load Statistics'.

Manage Jenkins


























-  [Configure System](#)
Configure global settings and paths.
-  [Configure Global Security](#)
Secure Jenkins; define who is allowed to access/use the system.
-  [Reload Configuration from Disk](#)
Discard all the loaded data in memory and reload everything from file system. Useful when you modified config files directly on disk.
-  [Manage Plugins](#)
Add, remove, disable or enable plugins that can extend the functionality of Jenkins. **(updates available)**
-  [System Information](#)
Displays various environmental information to assist trouble-shooting.
-  [System Log](#)
System log captures output from java.util.logging output related to Jenkins.
-  [Load Statistics](#)
Check your resource utilization and see if you need more computers for your builds.


Filter:


Updates Available **Installed** Advanced

Enabled	Name ↓	Version	Previously installed version	Pinned	Uninstall
<input checked="" type="checkbox"/>	Amazon EC2 plugin Allow Jenkins to start slaves on EC2 or Eucalyptus on demand, and kill them as they get unused.	1.29	Downgrade to 1.28		Uninstall
<input checked="" type="checkbox"/>	Ant Plugin This plugin adds Apache Ant support to Jenkins.	1.2			
<input checked="" type="checkbox"/>	build-env-propagator Copies environment variables added or modified during one build to the next.	1.0			Uninstall
<input checked="" type="checkbox"/>	CloudBees Build Flow plugin Manage jobs orchestration as a dedicated "build flow" top level item	0.18	Downgrade to 0.17		Uninstall
<input checked="" type="checkbox"/>	Cobertura Plugin This plugin integrates Cobertura coverage reports to Jenkins.	1.9.7	Downgrade to 1.9.5		Uninstall
<input checked="" type="checkbox"/>	Config File Provider Plugin Ability to provide configuration files (e.g. settings.xml for maven, XML, groovy, custom files,...) loaded through the UI which will be copied to the job workspace.	2.9.3			Uninstall
<input checked="" type="checkbox"/>	Copy Artifact Plugin Adds a build step to copy artifacts from another project.	1.35.2	Downgrade to 1.32.1		Uninstall
<input checked="" type="checkbox"/>	Credentials Plugin This plugin allows you to store credentials in Jenkins.	1.23	Downgrade to 1.18	Unpin ?	
<input checked="" type="checkbox"/>	CVS Plug-in Integrates Jenkins with CVS version control system using a modified version of the Netbeans cvsclient.	2.11	Downgrade to 2.8		
<input checked="" type="checkbox"/>	Dashboard View Customizable dashboard that can present various views of job information.	2.9.6	Downgrade to 2.9.4		Uninstall
<input checked="" type="checkbox"/>	disk-usage plugin This plugin counts disk usage.	0.25	Downgrade to 0.24		Uninstall
<input checked="" type="checkbox"/>	Durable Task Plugin Library offering an extension point for processes which can run outside of Jenkins yet be monitored.	1.6			Uninstall

Jenkins > Update center [ENABLE AUTO REFRESH](#)

Pipeline: Model API	 Success
Pipeline: Stage Tags Metadata	 Success
Pipeline: Declarative Agent API	 Success
Pipeline: Model Definition	 Success
SSH Credentials Plugin	 Success
Git client plugin	 Success
GIT server Plugin	 Success
Pipeline: Shared Groovy Libraries	 Success
Pipeline	 Success
GitHub API Plugin	 Success
Git plugin	 Success
GitHub plugin	 Success
GitHub Branch Source Plugin	 Success
GitHub Organization Folder Plugin	 Success
Pipeline: Stage View Plugin	 Success
Git plugin	 Success
MapDB API Plugin	 Success
Subversion Plug-in	 Success
SSH Slaves plugin	 Success
Matrix Authorization Strategy Plugin	 Success
PAM Authentication plugin	 Success
LDAP Plugin	 Success
Email Extension Plugin	 Success
Mailer Plugin	 Success
Kubernetes plugin	 Success






 [Go back to the top page](#)
(you can start using the installed plugins right away)

 Restart Jenkins when installation is complete and no jobs are running



 [Back to credential domains](#)

 **Add Credentials**

Kind	<input type="text" value="Username with password"/>	▼
Scope	<input type="text" value="Global (Jenkins, nodes, items, all child items, etc)"/>	▼ 
Username	<input type="text" value="admin"/>	
Password	<input type="password" value="....."/>	
ID	<input type="text"/>	
Description	<input type="text" value="Kubernetes Admin Credentials"/>	

OK

Cloud

Kubernetes

Name ?

Kubernetes URL ?

Kubernetes server certificate key

Disable https certificate check ?

Kubernetes Namespace

Credentials ?

Connection successful

Jenkins URL ?

Jenkins tunnel ?


Connection Timeout ?

Read Timeout ?

Container Cap ?

Images

List of Images to be launched as slaves

Kubernetes Namespace	<input type="text" value="default"/>
Credentials	<input type="text" value="admin/***** (Kubernetes Cluster Login)"/>  Add
	<input type="button" value="Test Connection"/>
Jenkins URL	<input type="text"/>
Jenkins tunnel	<input type="text"/>
Connection Timeout	<input type="text" value="5"/>
Read Timeout	<input type="text" value="15"/>
Container Cap	<input type="text" value="10"/>

Kubernetes Pod Template

Name

Labels

The name of the pod template to inherit from

Containers

Name

Docker image



Always pull image

Working directory



Command to run slave agent



Arguments to pass to the command



Allocate pseudo-TTY

EnvVars

Add

Environment



Variable

List of environment variables
to set in slave pod

Advanced...

Add

Cluster > Namespaces

Cluster





- Namespaces
- Nodes
- Persistent Volumes
- Roles
- Storage Classes

Namespace

default

Overview

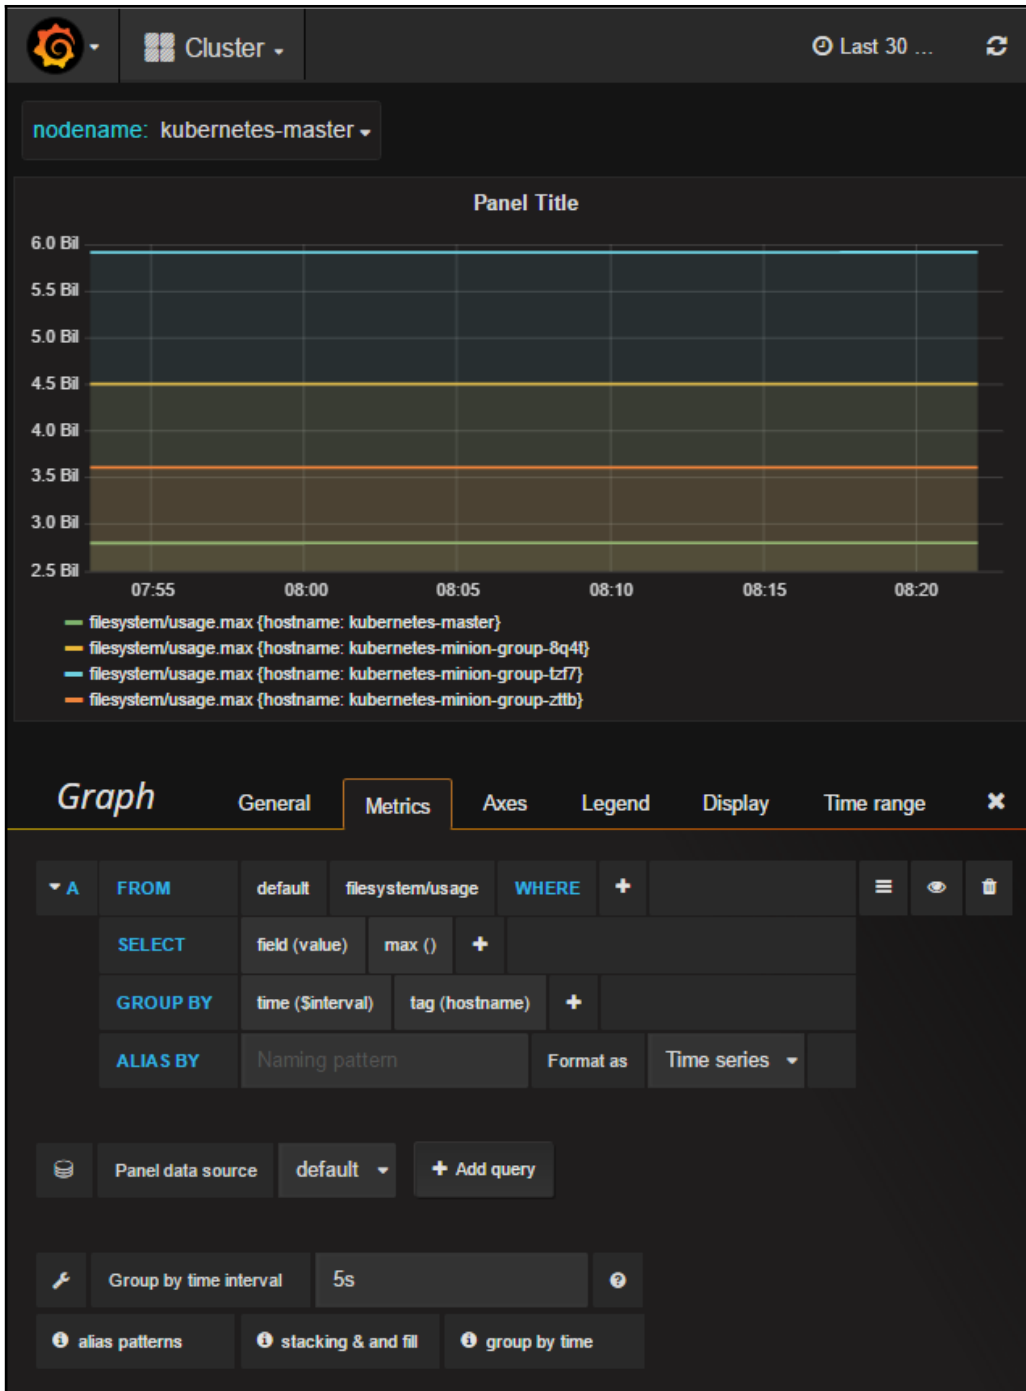
Namespaces

Name	Labels	Status	Age
 gsw-k8s-jenkins	-	Active	a day
 default	-	Active	5 days
 kube-public	-	Active	5 days
 kube-system	-	Active	5 days

Chapter 8: Monitoring and Logging

NAME	READY	STATUS	RESTARTS	AGE
etcd-empty-dir-cleanup-kubernetes-master	1/1	Running	2	2d
etcd-server-events-kubernetes-master	1/1	Running	2	2d
etcd-server-kubernetes-master	1/1	Running	2	2d
fluentd-cloud-logging-kubernetes-master	1/1	Running	2	2d
fluentd-cloud-logging-kubernetes-minion-group-rh7t	1/1	Running	0	3m
fluentd-cloud-logging-kubernetes-minion-group-s345	1/1	Running	0	3m
fluentd-cloud-logging-kubernetes-minion-group-tp2h	1/1	Running	0	3m
heapster-v1.2.0-2805816975-80mjc	4/4	Running	0	20h
kube-addon-manager-kubernetes-master	1/1	Running	2	2d
kube-apiserver-kubernetes-master	1/1	Running	4	2d
kube-controller-manager-kubernetes-master	1/1	Running	2	2d
kube-dns-4101612645-bwsd4	4/4	Running	0	20h
kube-dns-autoscaler-2715466192-gt3r7	1/1	Running	0	20h
kube-proxy-kubernetes-minion-group-rh7t	1/1	Running	0	4m
kube-proxy-kubernetes-minion-group-s345	1/1	Running	0	4m
kube-proxy-kubernetes-minion-group-tp2h	1/1	Running	0	3m
kube-scheduler-kubernetes-master	1/1	Running	2	2d
kubernetes-dashboard-3543765157-65g1m	1/1	Running	0	20h
l7-default-backend-2234341178-g4wct	1/1	Running	0	20h
l7-lb-controller-v0.8.0-kubernetes-master	1/1	Running	2	2d
monitoring-influxdb-grafana-v4-7x0n0	2/2	Running	0	20h
node-problem-detector-v0.1-lzfm1	1/1	Running	0	4m
node-problem-detector-v0.1-cjrtz	1/1	Running	0	4m
node-problem-detector-v0.1-f87pp	1/1	Running	2	2d
node-problem-detector-v0.1-vj001	1/1	Running	0	4m
rescheduler-v0.2.1-kubernetes-master	1/1	Running	2	2d

```
Name:          heapster-v1.2.0-2805816975-80mjc
Namespace:     kube-system
Node:          kubernetes-minion-group-rh7t/10.128.0.4
Start Time:    Thu, 02 Feb 2017 19:56:58 +0000
Labels:        k8s-app=heapster
                pod-template-hash=2805816975
                version=v1.2.0
Status:        running
IP:            10.244.7.5
Controllers:   ReplicaSet/heapster-v1.2.0-2805816975
```





Big value	Prefix		Value	avg	Postfix	
Font size	Prefix	50%	Value	80%	Postfix	50%
Unit	bytes			Decimals	auto	

Coloring

Background Value Thresholds ⓘ 184907710, 190907710

Colors invert order

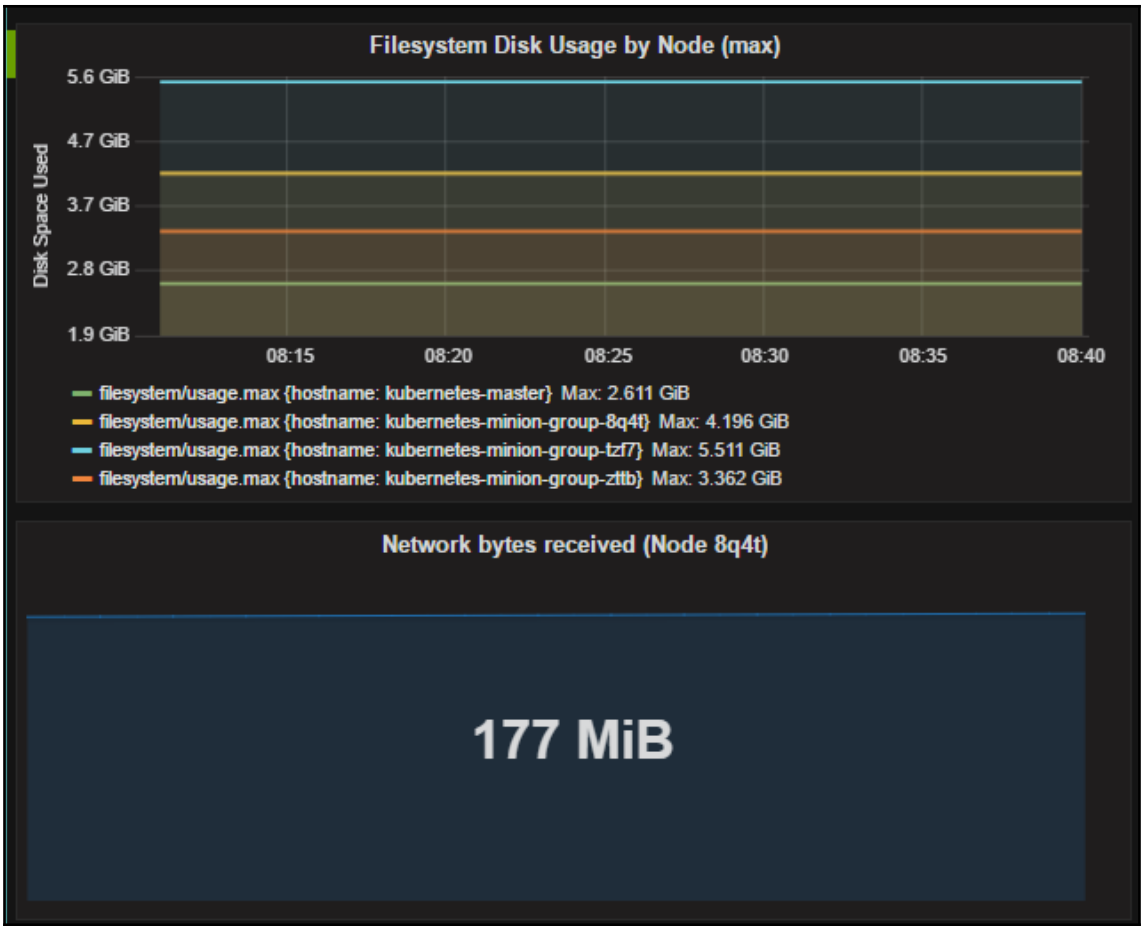
Spark lines

Show Background mode Line Color Fill Color

Gauge

Show Min 0 Max 100

Threshold labels Threshold markers



My First Project

CREATE METRIC CREATE EXPORT

Filter by label or text search

GCE VM Instance, kubernetes-minion-group... kubelet Any log level

Jump to date

2017-02-06 EST View Options

- ▶ i 11:12:47.000 MountVolume.Setup succeeded for volume "kubernetes.io/secret/9a2.
- ▶ i 11:12:48.000 GET /healthz: (39.901µs) 200 [[curl/7.26.0] 127.0.0.1:51762]
- ▶ i 11:12:50.000 MountVolume.Setup succeeded for volume "kubernetes.io/secret/9a0.
- ▶ i 11:12:53.000 MountVolume.Setup succeeded for volume "kubernetes.io/secret/99d.
- ▶ i 11:12:58.000 GET /healthz: (38.1µs) 200 [[curl/7.26.0] 127.0.0.1:51773]
- ▶ i 11:13:02.000 Found 44 PIDs in root, 44 of them are not to be moved
- ▶ i 11:13:05.000 GET /stats/summary/: (4.472253ms) 200 [[Go-http-client/1.1] 10.2.
- ▶ i 11:13:08.000 GET /healthz: (35.998µs) 200 [[curl/7.26.0] 127.0.0.1:51788]

Target

RESOURCE TYPE

Instance (GCE) ▼

APPLIES TO

Group ▼

kubernetes ▼

CONDITION TRIGGERS IF

Any Member Violates ▼

Configuration

IF METRIC

CPU Usage (GCE Monitoring) ▼

CONDITION


above ▼

THRESHOLD

80 %

FOR

5 minutes ▼



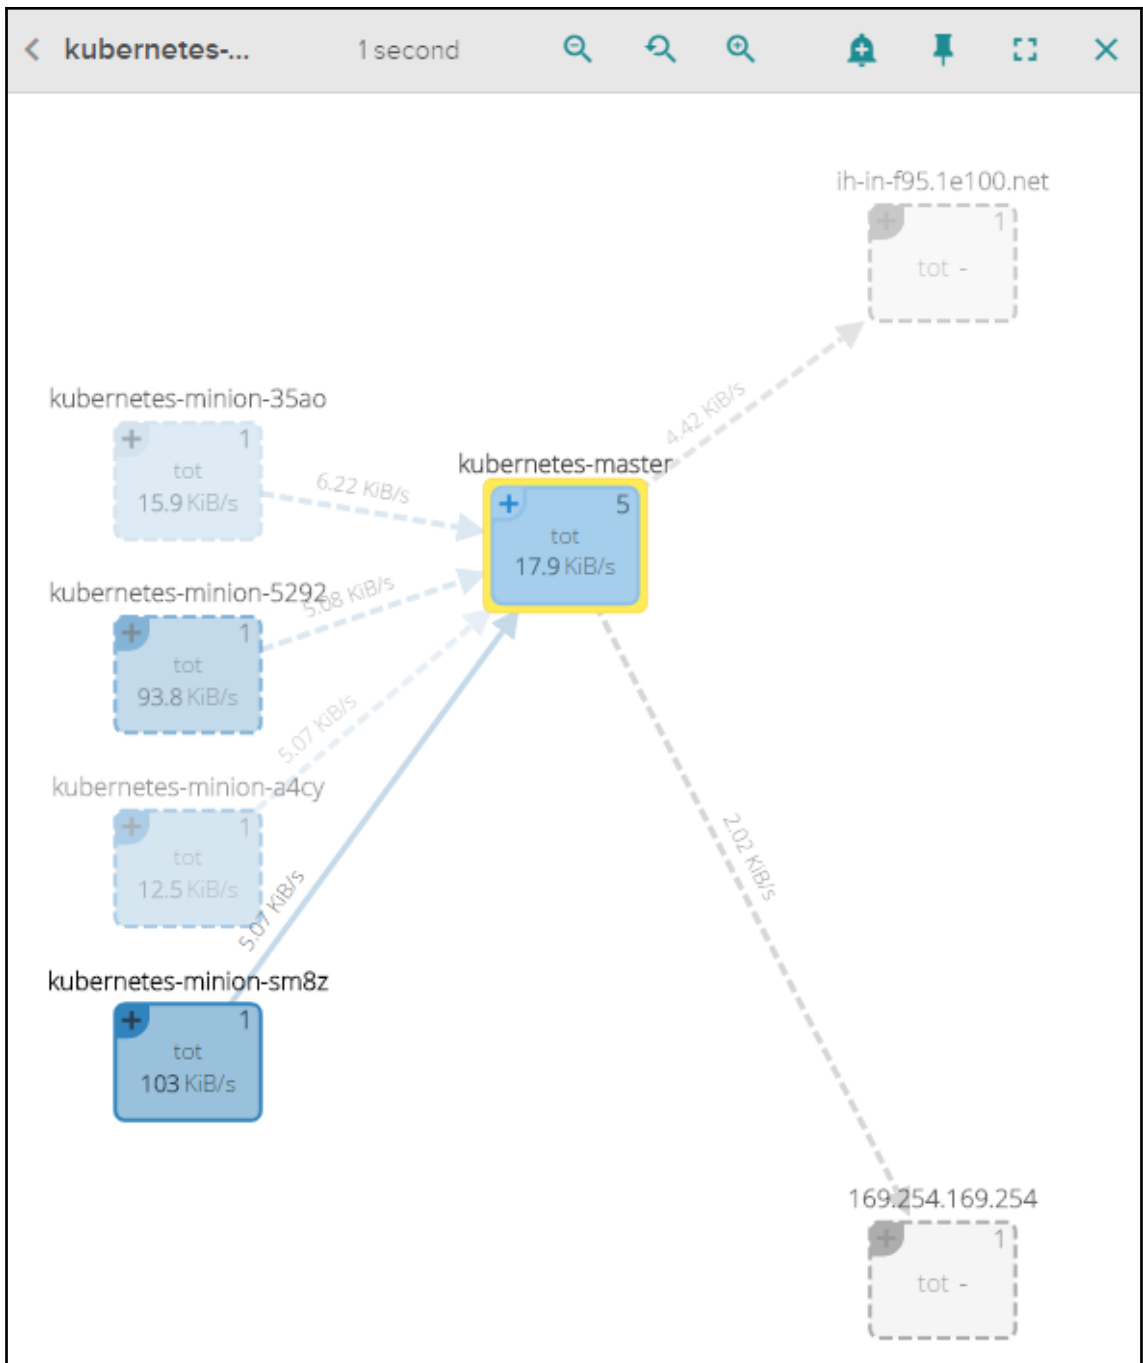
Your free trial will **expire in 14 days**, [upgrade your plan now!](#)

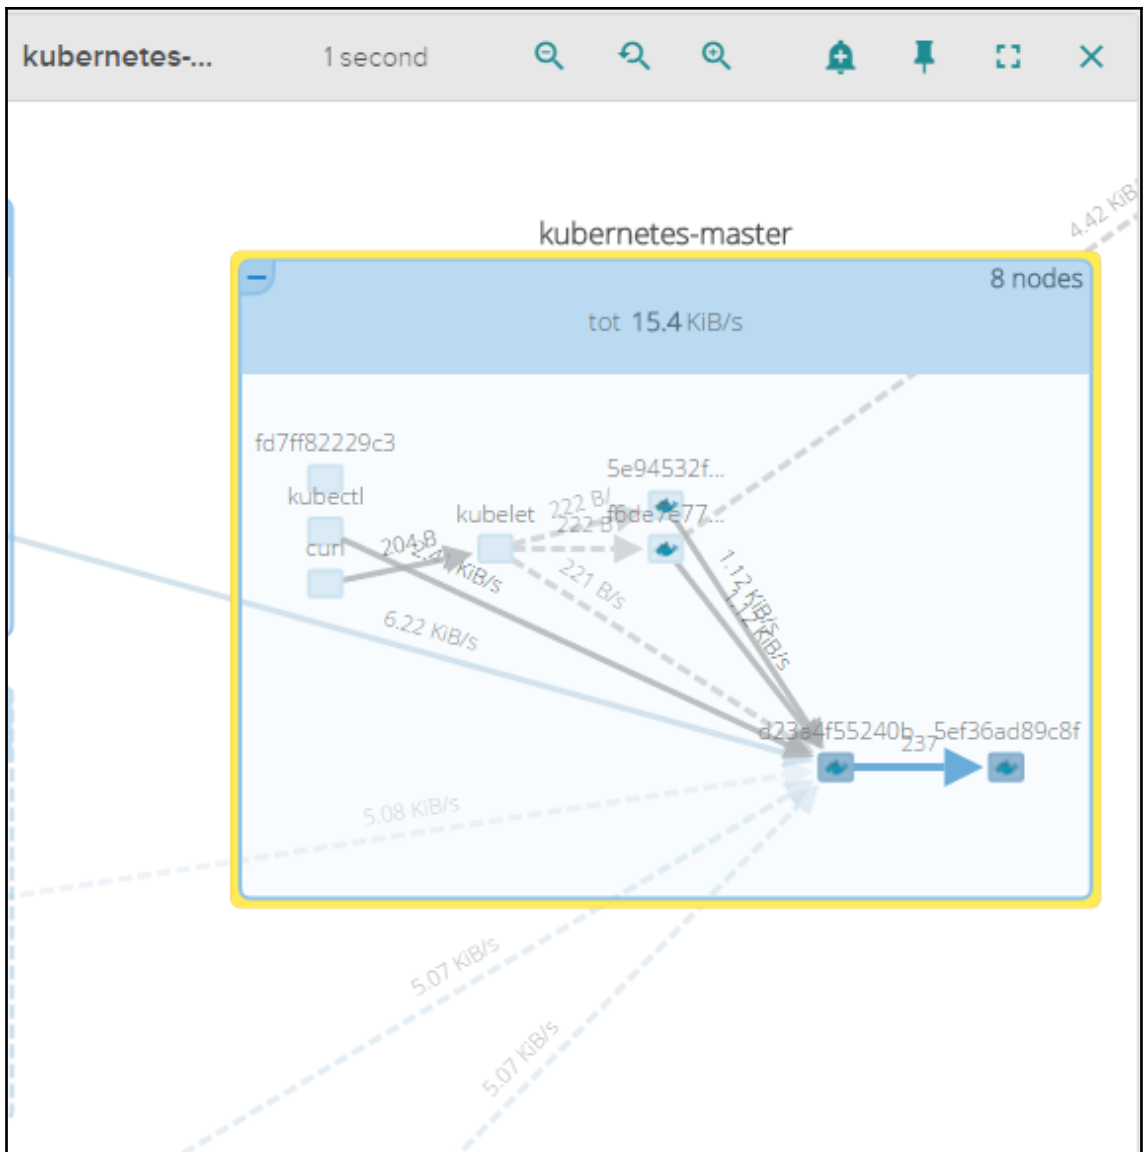
Events Alerts Captures J jon@responscomm.com

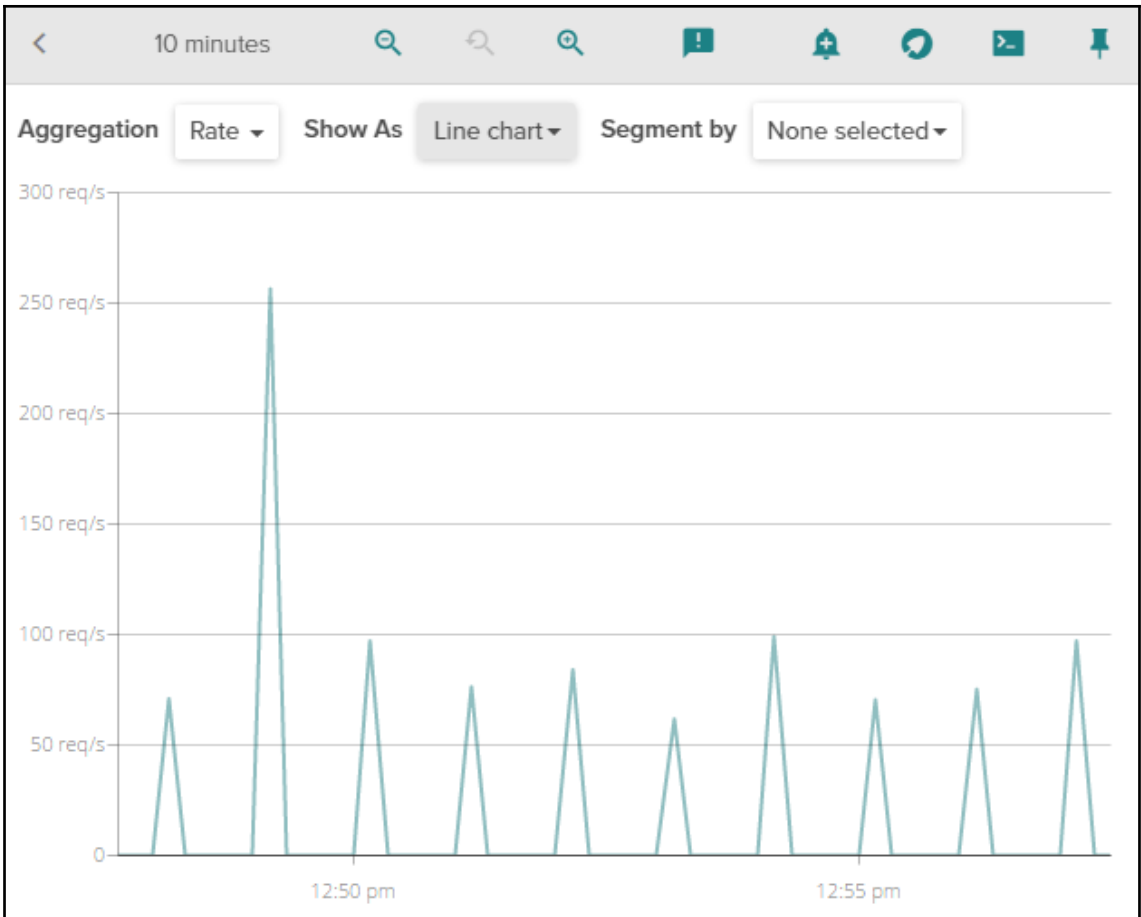
LIVE: LAST 10 SECONDS 10 S 1M 10 M 1H 6 H 1D 3 D 2 W

Overview Show host.m... contain... 10 s

Name	Instance Type	CPU %	Memory %	Network ... KiB/s
+ kubernetes-mini...	-	3.9	8.5	18.8
+ kubernetes-mini...	-	5.1	16.2	48.9







Name alert

×

Insert alert description

warning

1 Define Condition

MANUAL	BASELINE	HOST COMPARISON
<input checked="" type="checkbox"/> <code>cpu.idle.percent</code> ?	<input checked="" type="checkbox"/> <code>cpu.iowait.percent</code> ?	<input checked="" type="checkbox"/> <code>cpu.nice.percent</code> ?
<input checked="" type="checkbox"/> <code>cpu.stolen.percent</code> ?	<input checked="" type="checkbox"/> <code>cpu.system.percent</code> ?	<input checked="" type="checkbox"/> <code>cpu.used.percent</code> ?
<input checked="" type="checkbox"/> <code>cpu.user.percent</code> ?	<input checked="" type="checkbox"/> <code>file.bytes.total</code> ?	<input checked="" type="checkbox"/> <code>fs.used.percent</code> ?
<input checked="" type="checkbox"/> <code>memory.bytes.used</code> ?	<input checked="" type="checkbox"/> <code>memory.swap.bytes.available</code> ?	<input checked="" type="checkbox"/> <code>memory.swap.bytes.total</code> ?
<input checked="" type="checkbox"/> <code>memory.swap.bytes.used</code> ?	<input checked="" type="checkbox"/> <code>memory.swap.used.percent</code> ?	<input checked="" type="checkbox"/> <code>net.bytes.total</code> ?
<input checked="" type="checkbox"/> <code>net.request.count.in</code> ?	<input checked="" type="checkbox"/> <code>net.request.time.in</code> ?	<input checked="" type="checkbox"/> <code>net.tcp.queue.len</code> ?

Aggregated across everywhere ?

Segmented by none ?

2 Set Notifications

Configure your notification channels [here](#). Each alert you create can route notifications to any combination of these channels.

Email to `jon@responscomm.com` (Globally disabled)

This notification channel is globally disabled. Go to [Notification Settings page](#) to review settings.

3 Activate Sysdig Capture

Sysdig Capture lets you analyze a record of every system call executed on a host to troubleshoot containers, even after they have been deleted.

CANCEL

CREATE

Bytes	Process	Host_pid	Container_pid	container.name
79.06KB	kube-apiserver	5152	15	host
58.10KB	etcd	5211	10	host
6.29KB	dragent	19284	19292	host
4.52KB	kube-contr	5164	11	host
4.11KB	etcd	5211	11	host
1.95KB	kube-sched	5227	13	host
1.72KB	sshd	18963	18963	host

```

Viewing: Threads For: whole machine
Source: Live System Filter: evt.type!=switch

```

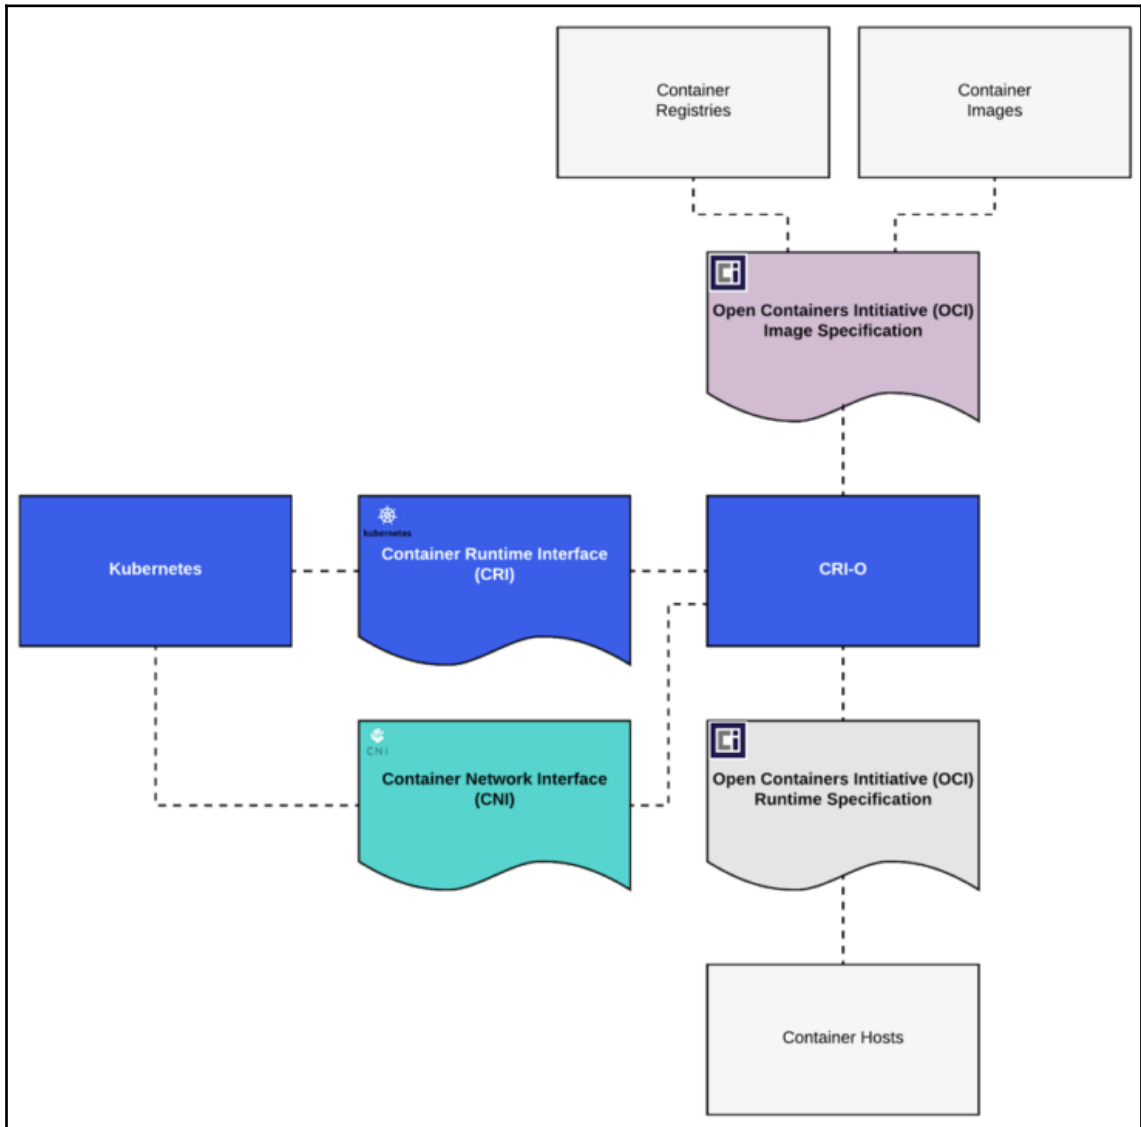
PID	TID	CPU	FILE	NET	Command
5152	5152	1.00	560.50	63.68K	/usr/local/bin/kube-apiserver --address=127.0.0.
5152	5153	0.50	0.00	0.00	/usr/local/bin/kube-apiserver --address=127.0.0.
5152	5154	0.00	0.00	0.00	/usr/local/bin/kube-apiserver --address=127.0.0.
5152	28713	0.00	0.00	0.00	/usr/local/bin/kube-apiserver --address=127.0.0.
5152	5161	0.00	0.00	0.00	/usr/local/bin/kube-apiserver --address=127.0.0.
5152	13254	0.00	191.00	6.75K	/usr/local/bin/kube-apiserver --address=127.0.0.
5152	15161	0.00	100.50	2.50	/usr/local/bin/kube-apiserver --address=127.0.0.

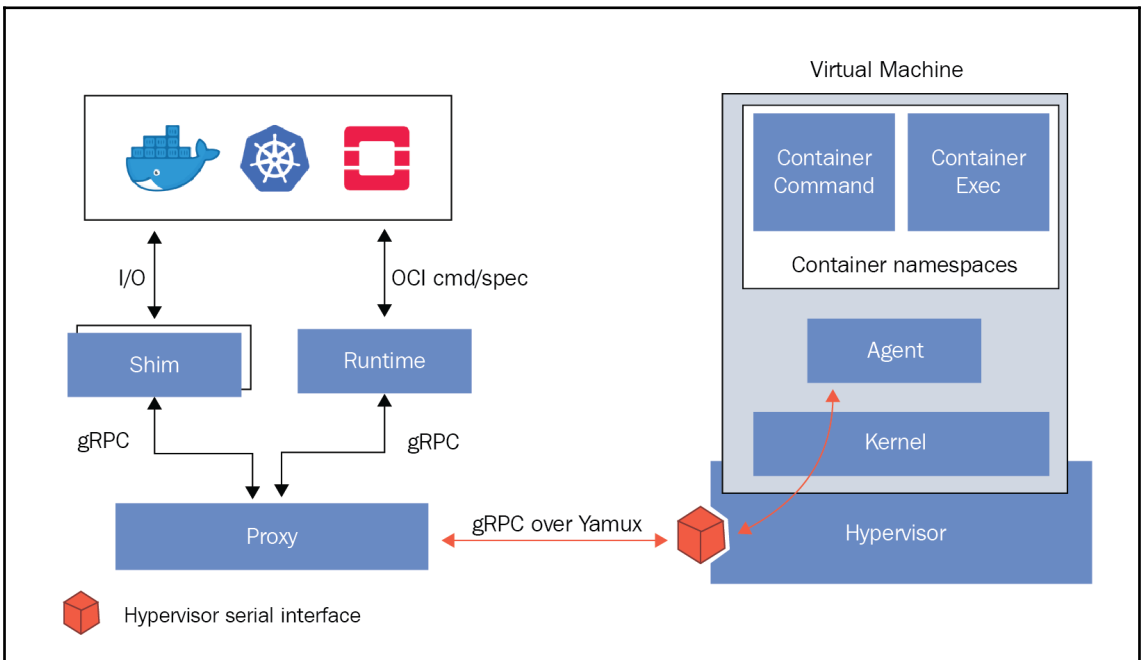
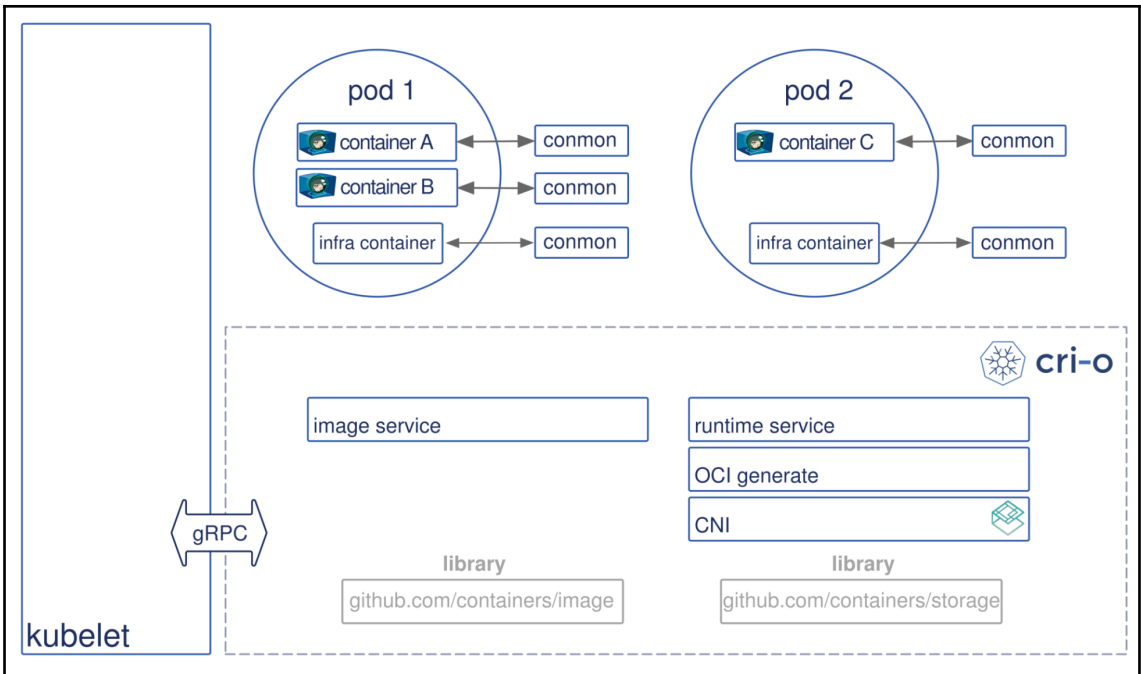
```

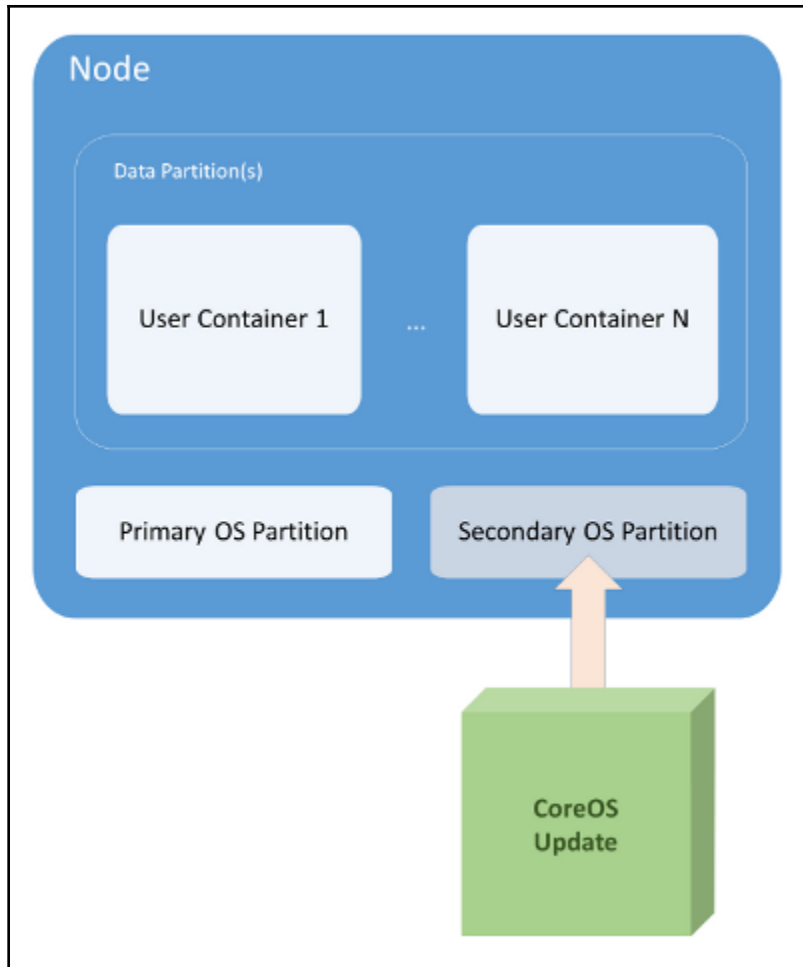
F1 Help F2 Views F4 Filter F5 Echo F6 Dig F7 Legend CTRL+F Search p Pause 1/7(14.3%)

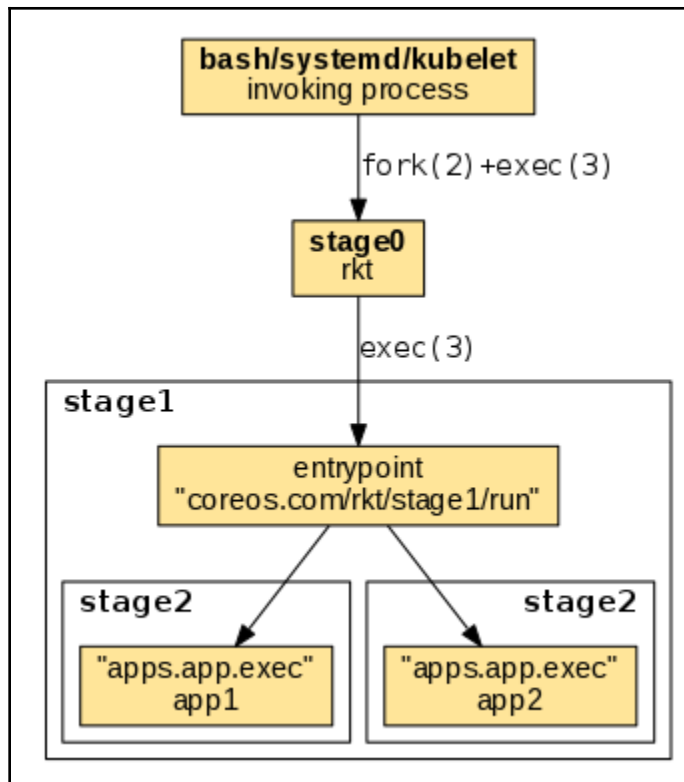
```


Chapter 9: Operating Systems, Platforms, and Cloud and Local Providers









 Browse Cluster

Deployments

Services

Jobs

Replica Sets

Daemon Sets

Replication Controllers

Autoscalers

Pods

Service Accounts


Config Maps

Secrets

Events

Search

Ingress

 Administration

Namespaces

Nodes

Cluster Status

Cluster Health

Tectonic Console	✓ All systems go
------------------	------------------

Kubernetes API Connection	✓ All systems go
---------------------------	------------------


Software Details

Kubernetes	v1.5.2+coreos.1
------------	-----------------

Tectonic	1.5.2-tectonic.1 Release Notes
----------	---

License	10 Nodes
---------	----------

Cloud Provider	Amazon Web Services
----------------	---------------------



Namespace: all ▾

RC node-js

Overview Edit Pods Events

Replication Controller Details

Replicas:
- +

Controller Labels:
deployment=demo x name=node-js x
app=frontend

Labels for this controller.

Label Query:
deployment=demo x name=node-js x
app=frontend

Write a label query that will match labels on new or existing pods.

Desired Pod State

These containers make up a pod. All of these containers are deployed together onto nodes in the cluster.

Pod Labels:

- Browse Cluster
- Deployments
- Services
- Jobs
- Replica Sets
- Daemon Sets
- Replication Controllers**
- Autoscalers
- Pods
- Service Accounts
- Config Maps
- Secrets
- Events
- Search
- Ingress
- Administration
 - Namespaces
 - Nodes
 - Cluster Settings
 - Roles
 - Role Bindings
 - Cluster Roles

Cluster Role Bindings

admin

My Account

Log Out

Pod Labels:

deployment=demo x name=node-js x
app=frontend

Each pod instance will have these labels. Services matching these labels will automatically send traffic to containers.

Containers

Add Another Container

CONTAINER NAME
node-js

CONTAINER IMAGE
jonbaier/node-express-info

CONTAINER VERSION/TAG
latest


PORTS
0 Ports >

PRIMARY COMMAND
Default Command >

PULL POLICY
Always Pull >

Update Replication Controller Cancel

[76]





Namespace: all ▾

- Browse Cluster
- Deployments
- Services
- Jobs
- Replica Sets
- Daemon Sets
- Replication Controllers
- Autoscalers
- Pods
- Service Accounts
- Config Maps
- Secrets
- Events
- Search
- Ingress
- ⚙ Administration
- Namespaces
- Nodes
- Cluster Settings
- Roles
- Role Bindings
- Cluster Roles
- Cluster Role Bindings
- 👤 admin
- My Account
- Log Out


Events

All Types ▾
All Categories ▾


 Streaming events...
Showing 245 events

 node-js-lskv6


Successfully pulled image "jonbaier/node-express-info:latest"

 7 minutes ago


Generated from kubelet on [ip-10-0-116-242.us-west-1.compute.internal](#)

 node-js-lskv6


Created container with docker id c2e32858c8e6; Security: [seccomp=unconfined]

 7 minutes ago


Generated from kubelet on [ip-10-0-116-242.us-west-1.compute.internal](#)

 node-js-lskv6

Started container with docker id c2e32858c8e6

 7 minutes ago

Generated from kubelet on [ip-10-0-116-242.us-west-1.compute.internal](#)

 node-js-86qcm

The screenshot shows the Tectonic dashboard interface. On the left is a dark blue sidebar with the Tectonic logo and a navigation menu. The main content area has a blue header with 'Namespace: all' and a dropdown menu open, listing 'all', 'default', 'kube-system', and 'tectonic-system'. Below the dropdown is a 'Filter Deployments by name...' input field. A table below shows a deployment named 'default-http-backend' with the label 'app=default-http-backend'.

TECTONIC
by CoreOS

Namespace: all ▾

Deployments


all
default
kube-system
tectonic-system

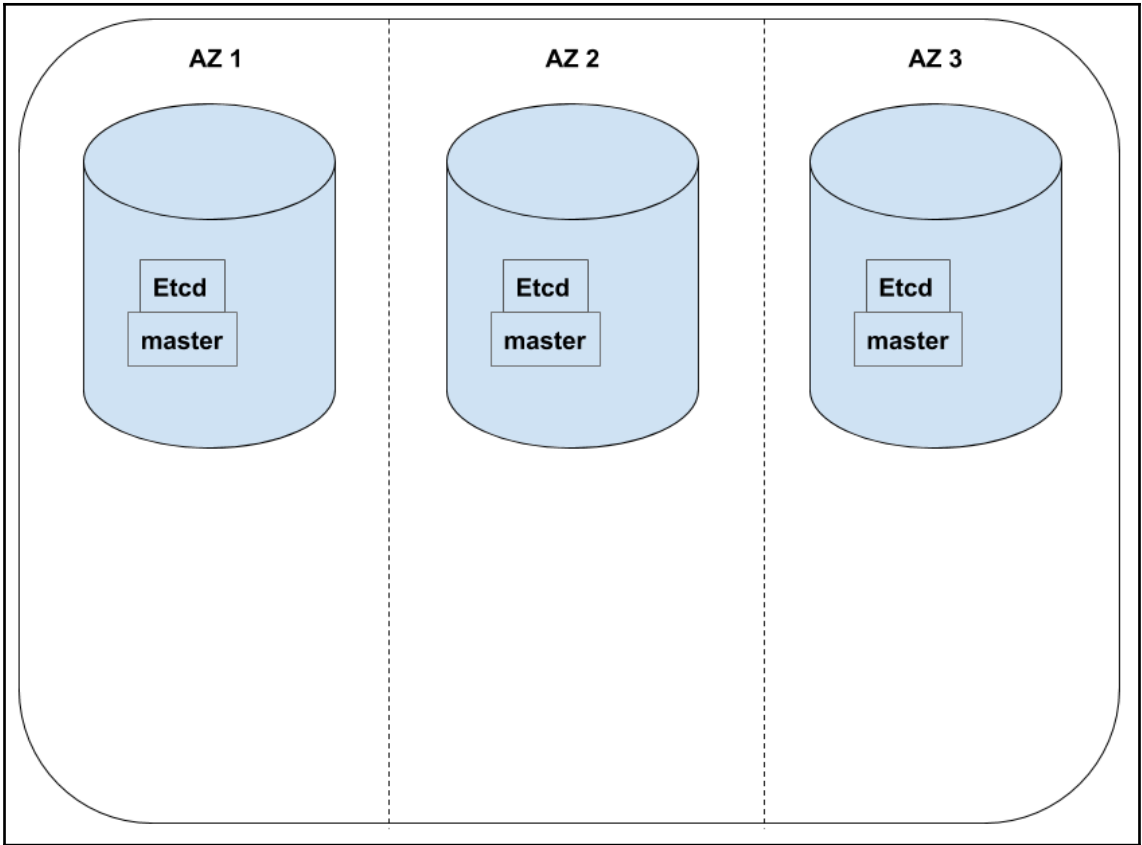
Create Deployment

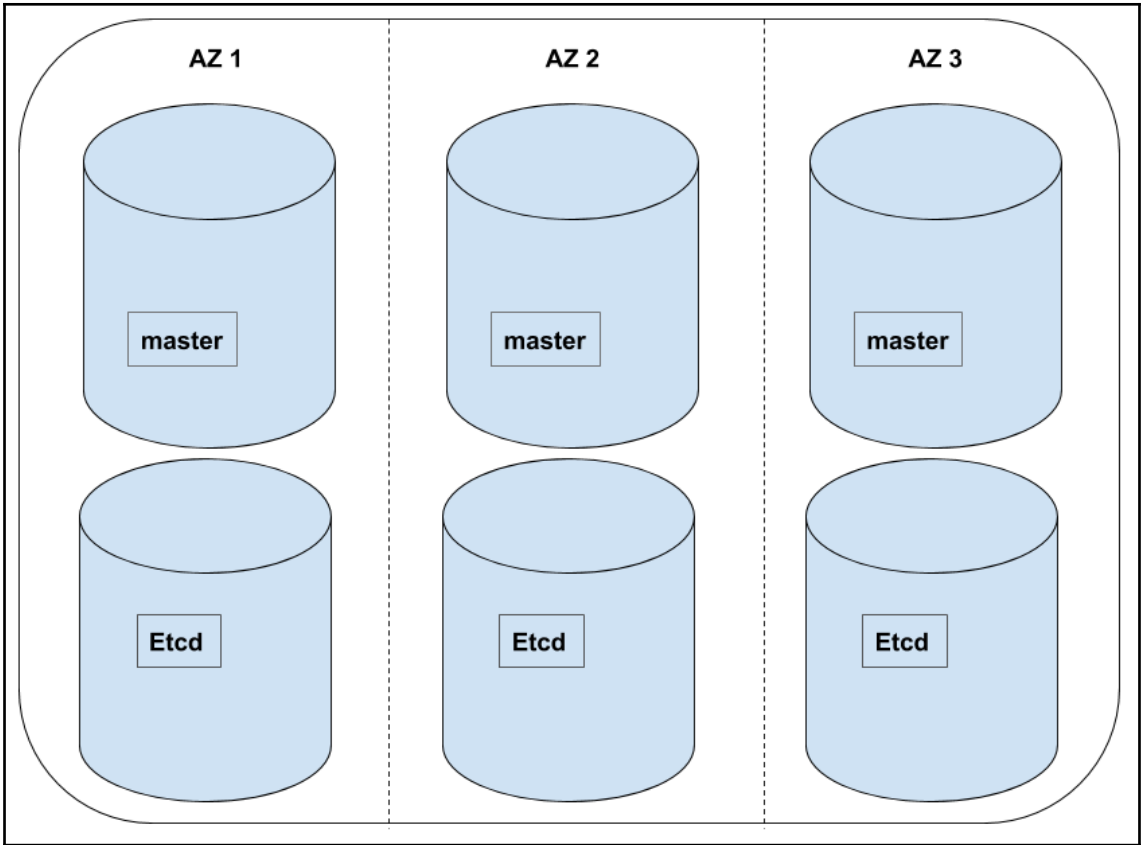
Filter Deployments by name...

NAME	LABELS
default-http-backend	app=default-http-backend

Chapter 10: Designing for High Availability and Scalability













kubernetes / community Watch 455 Star 2,424 Fork 1,416


Code Issues 104 Pull requests 113 Projects 0 Insights

Branch: master community / sig-apps / Create new file Upload files Find file History

 nikhita generator: remove google groups note from all SIG READMEs Latest commit 180425b 4 days ago

..

 minutes	> fix some typos	7 months ago
 CONTRIBUTING.md	Initial contribution of SIG Apps contributions guide	7 months ago
 OWNERS	Add sig labels and use aliases in OWNERS files	10 months ago
 README.md	generator: remove google groups note from all SIG READMEs	4 days ago
 agenda.md	Change zoom links per SC meeting	2 months ago

 README.md

Apps Special Interest Group

Covers deploying and operating applications in Kubernetes. We focus on the developer and devops experience of running applications in Kubernetes. We discuss how to define and run apps in Kubernetes, demo relevant tools and projects, and discuss areas of friction that can lead to suggesting improvements or feature requests.

Meetings

- Regular SIG Meeting: [Mondays at 9:00 PT \(Pacific Time\)](#) (weekly). [Convert to your timezone](#).
 - [Meeting notes and Agenda](#).
 - [Meeting recordings](#).

Leadership

Chairs

The Chairs of the SIG run operations and processes governing the SIG.

- Matt Farina ([@mattfarina](#)), Samsung SDS
- Adnan Abdulhussein ([@prydonius](#)), Bitnami
- Kenneth Owens ([@kow3ns](#)), Google

Contact

Consider API aggregation if:	Prefer a stand-alone API if:
Your API is Declarative .	Your API does not fit the Declarative model.
You want your new types to be readable and writable using kubectl .	kubectl support is not required
You want to view your new types in a Kubernetes UI, such as dashboard, alongside built-in types.	Kubernetes UI support is not required.
You are developing a new API.	You already have a program that serves your API and works well.
You are willing to accept the format restriction that Kubernetes puts on REST resource paths, such as API Groups and Namespaces. (See the API Overview .)	You need to have specific REST paths to be compatible with an already defined REST API.
Your resources are naturally scoped to a cluster or to namespaces of a cluster.	Cluster or namespace scoped resources are a poor fit; you need control over the specifics of resource paths.
You want to reuse Kubernetes API support features .	You don't need those features.

Google Cloud Platform GSW K8s 3

DASHBOARD ACTIVITY

Project info

- Project name: GSW K8s 3
- Project ID: gsw-k8s-3
- Project number: 803769040293

[Go to project settings](#)

Resources

This project has no resources

Trace

No trace data from the past 7 days

[Get started with Stackdriver Trace](#)

Getting Started

- API Enable APIs and get credentials like keys
- Deploy a prebuilt solution
- Add dynamic logging to a running application
- Monitor errors with Error Reporting
- Deploy a Hello World app
- Take a VM quickstart
- Create a Cloud Storage bucket
- Create a Cloud Function
- Install the Cloud SDK

[Explore all tutorials](#)

API APIs

Requests (requests/sec)

[Go to APIs overview](#)

Google Cloud Platform status

All services normal

[Go to Cloud status dashboard](#)

Billing

Estimated charges USD
For the billing period Sep 1 - 3, 2018

[View detailed charges](#)

Error Reporting

No sign of any errors. Have you set up Error

[Learn how to set up Error Reporting](#)

News

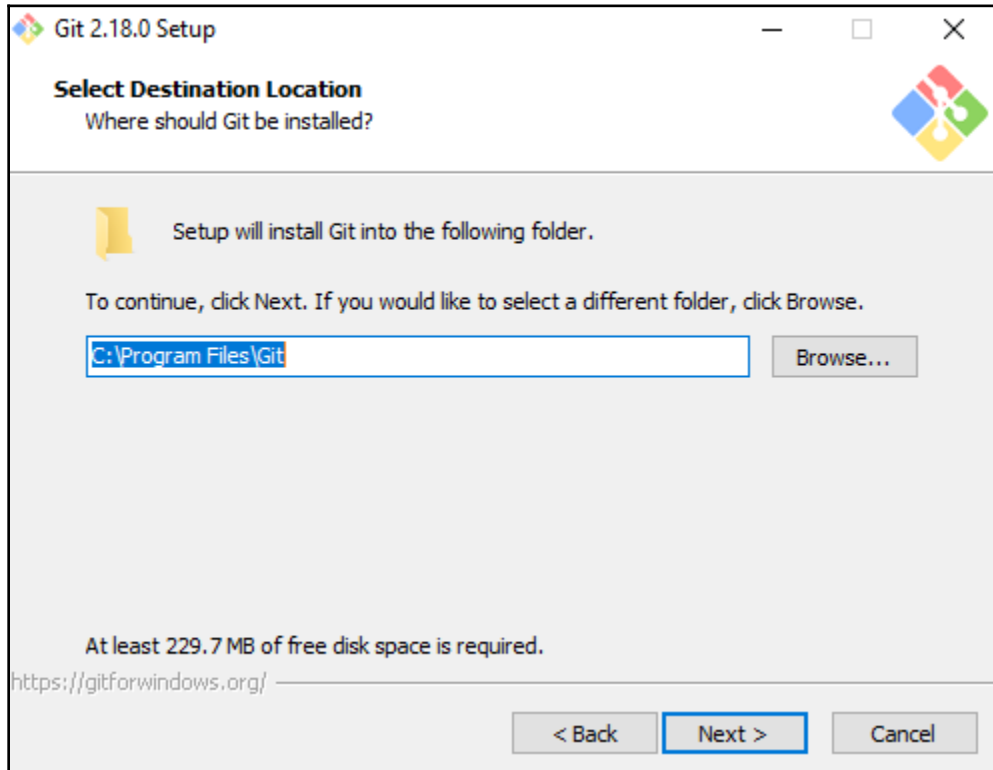
- We've moved! Come see our new home! Aug 6, 2018
- Last month today: July on GCP Aug 3, 2018
- Repairing network hardware at scale with SRE Aug 1, 2018

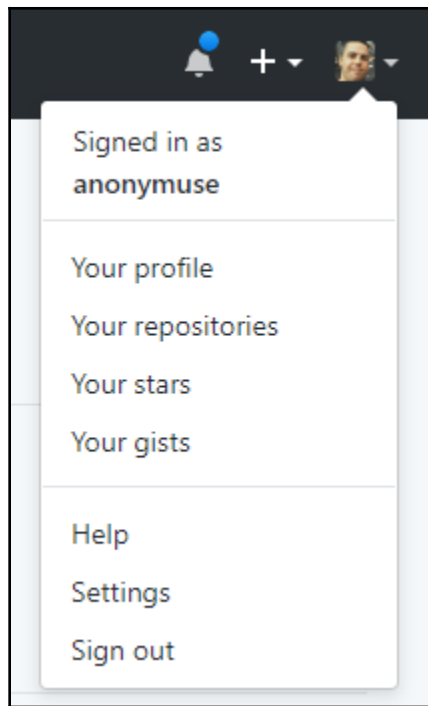
[Read all news](#)

Documentation

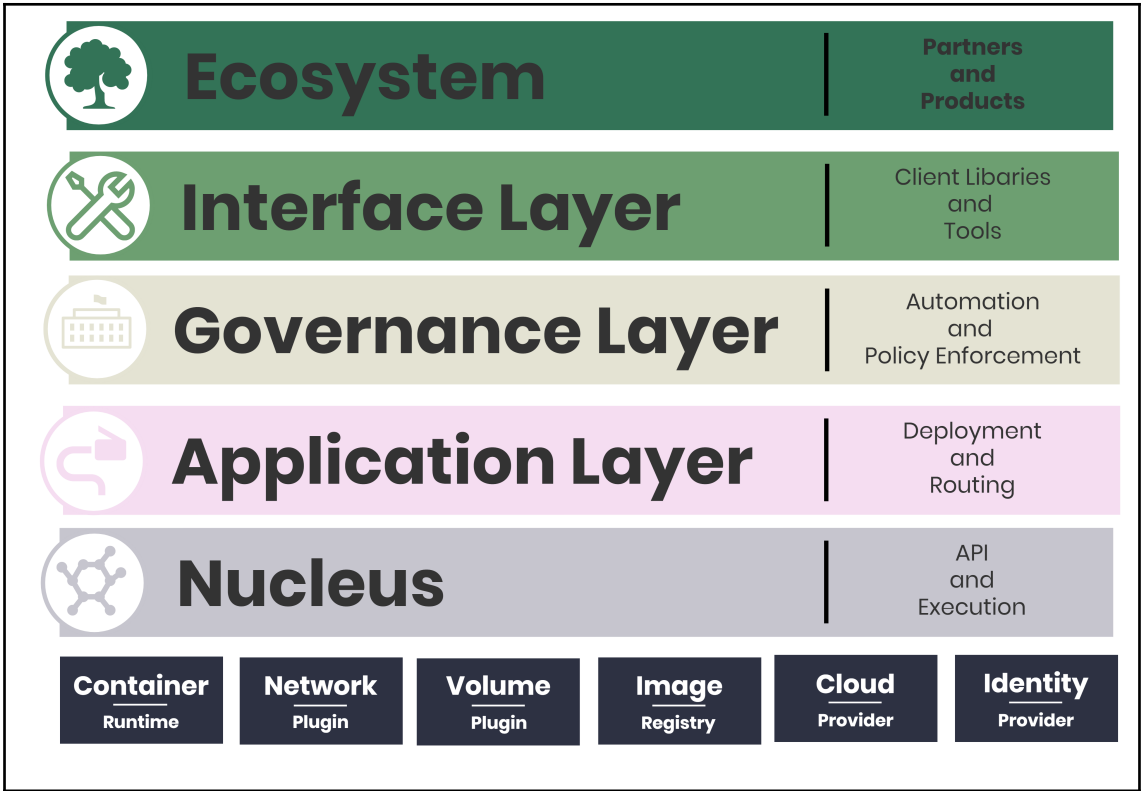
- Learn about Compute Engine
- Learn about Cloud Storage
- Learn about App Engine

Chapter 11: Kubernetes SIGs, Incubation Projects, and the CNCF





Personal settings
Profile
Account
Emails
Notifications
Billing
SSH and GPG keys
Security
Blocked users
Repositories
Organizations
Saved replies
Applications





CLOUD NATIVE TRAIL MAP

The Cloud Native Landscape l.cncf.io has a large number of options. This Cloud Native Trail Map is a recommended process for leveraging open source, cloud native technologies. At each step, you can choose a vendor-supported offering or do it yourself, and everything after step #3 is optional based on your circumstances.

HELP ALONG THE WAY

A. Training and Certification

Consider training offerings from CNCF and then take the exam to become a Certified Kubernetes Administrator or a Certified Kubernetes Application Developer cncf.io/training

B. Consulting Help

If you want assistance with Kubernetes and the surrounding ecosystem, consider leveraging a Kubernetes Certified Service Provider

cncf.io/kcsp

C. Join CNCF's End User Community

For companies that don't offer cloud native services externally cncf.io/enduser

WHAT IS CLOUD NATIVE?

Cloud-native technologies, such as containers and microservices, empower organizations to develop and deploy scalable, agile applications and services in dynamic, distributed environments. By taking into account these characteristics, such systems are designed to be resilient, elastic, and loosely coupled, via manageable abstractions and declarative APIs, thereby enabling effective, reliable automation. This allows engineers to observe the applications and to safely make impactful changes, and results in processes and workflows that fully take advantage of these environments and minimize toil.

The Cloud Native Computing Foundation seeks to drive adoption of these techniques by fostering an ecosystem of open-source, vendor-neutral projects that align with these objectives, and which are portable to public, private, and hybrid clouds. We democratize the state-of-the-art patterns and practices to ensure innovations remain open and accessible for everyone.

l.cncf.io

v20180604



1. CONTAINERIZATION

- Commonly done with Docker containers
- Any size application and dependencies (even PDP-11 code running on an emulator) can be containerized
- Over time, you should aspire towards splitting suitable applications and writing future functionality as microservices



3. ORCHESTRATION & APPLICATION DEFINITION

- Kubernetes is the market-leading orchestration solution
- You should select a Certified Kubernetes Distribution, Hosted Platform, or Installer: cncf.io/ck
- Helm Charts help you define, install, and upgrade even the most complex Kubernetes application



5. SERVICE MESH AND DISCOVERY

- CoreDNS is a fast and flexible tool that is useful for service discovery
- Envoy and Linkerd each enable service mesh architectures
- They offer health checking, routing, and load balancing



7. DISTRIBUTED DATABASE

When you need more resiliency and scalability than you can get from a single database, Vitess is a good option for running MySQL at scale through sharding.



9. CONTAINER RUNTIME

You can use alternative container runtimes. The most common, all of which are OCI-compliant, are containerd, rkt and CRI-O.



2. CI/CD

- Setup Continuous Integration/Continuous Delivery (CI/CD) so that changes to your source code automatically result in a new container being built, tested, and deployed to staging and eventually, perhaps, to production
- Setup automated rollouts, roll backs and testing



4. OBSERVABILITY & ANALYSIS

- Pick solutions for monitoring, logging and tracing
- Consider CNCF projects Prometheus for monitoring, Fluentd for logging and Jaeger for Tracing
- For tracing, look for an OpenTracing-compatible implementation like Jaeger



6. NETWORKING

To enable more flexible networking, use a CNI-compliant network project like Calico, Flannel, or Weave Net.



8. MESSAGING

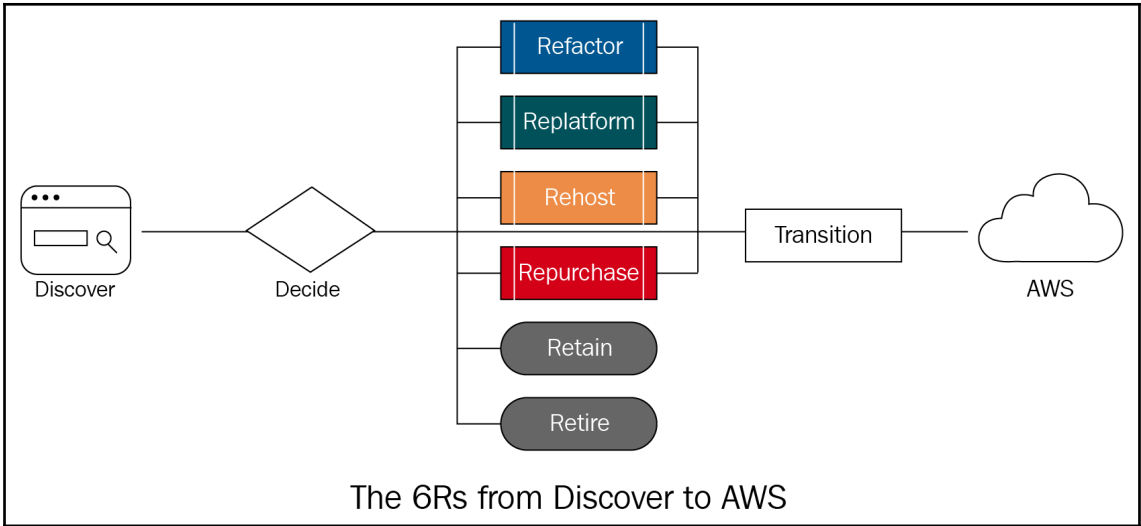
When you need higher performance than JSON-RPC, consider using gRPC. NATS is publish/subscribe message-oriented middleware.



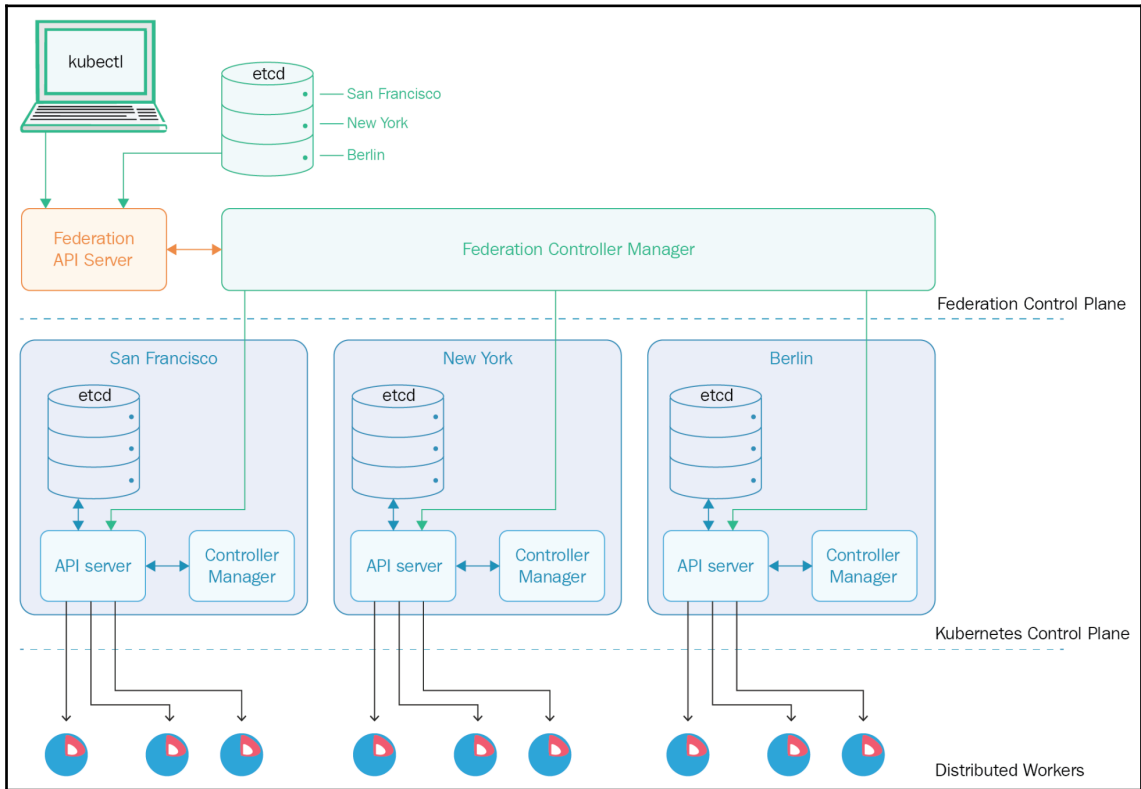
10. SOFTWARE DISTRIBUTION

If you need to do secure software distribution, evaluate Notary, an implementation of The Update Framework.





Chapter 12: Cluster Federation and Multi-Tenancy



CURRENT	NAME	CLUSTER	AUTHINFO	NAMESPACE
*	awsk8s	awsk8s	awsk8s	
	gcek8s	gcek8s	gcek8s	

CURRENT	NAME	CLUSTER	AUTHINFO	NAMESPACE
*	awsk8s	awsk8s	awsk8s	
	gcek8s	gcek8s	gcek8s	
	master-control	master-control	master-control	

NAME	READY	STATUS
RESTARTS AGE		
master-control-apiserver-3595964982-s6lx9 0 8m	2/2	Running
master-control-controller-manager-516854663-r8m37 0 8m	1/1	Running

the server doesn't have a resource type "pods"

```

Name:          node-js-deploy
Namespace:    default
CreationTimestamp:  Fri, 10 Mar 2017 22:15:11 +0000
Labels:       name=node-js-deploy
Selector:     name=node-js-deploy
Replicas:    0 updated | 3 total | 3 available | 0 unavailable
StrategyType: RollingUpdate
MinReadySeconds: 0
RollingUpdateStrategy: 1 max unavailable, 1 max surge
Events:
  FirstSeen      LastSeen      Count   From
  SubObjectPath  Type           Reason      Message
  -----
  4m             4m            1        {federated-deployment-controller }
Normal          CreateInCluster Creating deployment in cluster gcek8s
  4m             4m            1        {federated-deployment-controller }
Normal          CreateInCluster Creating deployment in cluster awsk8s

```



LASTSEEN	FIRSTSEEN	COUNT	NAME	KIND	SUBJECT	TYPE
REASON	SOURCE	MESSAGE				
10m	10m	1	node-js-deploy	Deployment	{federated-deployment-controller }	Normal
CreateInCluster		Creating deployment in cluster gcek8s				
10m	10m	1	node-js-deploy	Deployment	{federated-deployment-controller }	Normal
CreateInCluster		Creating deployment in cluster awsk8s				

NAME	READY	STATUS	RESTARTS	AGE
node-js-deploy-1713031517-1661z	1/1	Running	0	7m

NAME	READY	STATUS	RESTARTS	AGE
node-js-deploy-1713031517-bvdmf	1/1	Running	0	7m
node-js-deploy-1713031517-jnfnr	1/1	Running	0	7m

```
apiVersion: v1
data:
  backend-service.url: my-backend-service
kind: ConfigMap
metadata:
  creationTimestamp: 2017-03-10T22:28:38Z
  name: my-application-config
  namespace: default
  resourceVersion: "1959"
  selfLink: /api/v1/namespaces/default/configmaps/my-application-config
  uid: e85a0028-05e0-11e7-bdf8-42010a800002
```

NAME	READY	STATUS	RESTARTS	AGE
node-js-deploy-1713031517-cmd7q	1/1	Running	0	39m
node-js-deploy-1713031517-zncxr	1/1	Running	0	39m
node-js-rs-6g7nj	1/1	Running	0	9m
node-js-rs-f4w7b	1/1	Running	0	9m



```

Cloud Shell
(gsw-k8s-3) x (gsw-k8s-3) +
> --num-nodes "4" --network "default" --enable-cloud-logging --enable-cloud-monitoring --enable-ip-ali
WARNING: Starting in 1.12, new clusters will not have a client certificate issued. You can manually enab
ing the `--[no-]issue-client-certificate` flag.
This will enable the autorepair feature for nodes. Please see
https://cloud.google.com/kubernetes-engine/docs/node-auto-repair for more
information on node autorepairs.

WARNING: Starting in Kubernetes v1.10, new clusters will no longer get compute-rw and storage-ro scopes
will remain included in the default --scopes). To use these scopes, add them explicitly to --scopes. To
property (gcloud config set container/new_scopes_behavior true).
NAME          TYPE      LOCATION  TARGET  STATUS_MESSAGE  STATUS      START_TIME  END_TIME
cluster-1     us-east1-b  PROVISIONING
  
```

Google Cloud Platform GSW K8s 3

Kubernetes Engine

Kubernetes clusters

A Kubernetes cluster is a managed group of uniform VM instances for running Kubernetes. [Learn more](#)

Filter by label or name

<input type="checkbox"/> Name ^	Location	Cluster size	Total cores	Total memory	Notifications	Labels
<input type="checkbox"/> <input checked="" type="checkbox"/> cluster-1	us-east1-b	4	8 vCPUs	30.00 GB	Low resource requests	Connect
<input type="checkbox"/> <input checked="" type="checkbox"/> cluster-2	us-east1-b	4	8 vCPUs	30.00 GB	Low resource requests	Connect

Firewall rules

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. [Learn more](#)

Note: App Engine firewalls are managed [here](#).

Filter resources

<input type="checkbox"/> Name	Type	Targets	Filters	Protocols / ports	Action	Priority	Network ^
<input type="checkbox"/> istio-multicluster-test-pods	Ingress	gke-cluster-1-37037bd0-node, gke-cluster-2-909a776f-node	IP ranges: 10.40.0.0/14, 10.8.0.0/14	tcp udp icmp; esp; 2 more	Allow	900	default
<input type="checkbox"/> gke-cluster-1-37037bd0-all	Ingress	gke-cluster-1-37037bd0-node	IP ranges: 10.8.0.0/14	tcp udp sctp; icmp; 2 more	Allow	1000	default
<input type="checkbox"/> gke-cluster-1-37037bd0-ssh	Ingress	gke-cluster-1-37037bd0-node	IP ranges: 35.237.54.93/32	tcp:22	Allow	1000	default
<input type="checkbox"/> gke-cluster-1-37037bd0-vms	Ingress	gke-cluster-1-37037bd0-node	IP ranges: 10.128.0.0/9	tcp:1-65535 udp:1-65535 icmp	Allow	1000	default
<input type="checkbox"/> gke-cluster-2-909a776f-all	Ingress	gke-cluster-2-909a776f-node	IP ranges: 10.40.0.0/14	tcp udp icmp; esp; 2 more	Allow	1000	default
<input type="checkbox"/> gke-cluster-2-909a776f-ssh	Ingress	gke-cluster-2-909a776f-node	IP ranges: 35.237.47.212/32	tcp:22	Allow	1000	default
<input type="checkbox"/> gke-cluster-2-909a776f-vms	Ingress	gke-cluster-2-909a776f-node	IP ranges: 10.128.0.0/9	tcp:1-65535 udp:1-65535 icmp	Allow	1000	default
<input type="checkbox"/> default-allow-icmp	Ingress	Apply to all	IP ranges: 0.0.0.0/0	icmp	Allow	65534	default
<input type="checkbox"/> default-allow-internal	Ingress	Apply to all	IP ranges: 10.128.0.0/9	tcp:0-65535 udp:0-65535 icmp	Allow	65534	default
<input type="checkbox"/> default-allow-rdp	Ingress	Apply to all	IP ranges: 0.0.0.0/0	tcp:3389	Allow	65534	default
<input type="checkbox"/> default-allow-ssh	Ingress	Apply to all	IP ranges: 0.0.0.0/0	tcp:22	Allow	65534	default

Rows per page: 50 1 - 11 of 11

← Clusters
✎ EDIT
🗑️ DELETE
+ DEPLOY
🔗 CONNECT

⚠️ cluster-1

ℹ️ Deleting cluster.
The values shown below are going to change soon.

Details
Storage
Nodes

Cluster

Master version	1.10.6-gke.2
Endpoint	35.237.54.93 Show credentials
Client certificate	Enabled
Binary authorization	Disabled
Kubernetes alpha features	Disabled
Total size	4
Master zone	us-east1-b
Node zones	us-east1-b
Network	default
Subnet	default
VPC-native (alias IP)	Enabled
Pod address range	10.8.0.0/14
Service address range	10.12.0.0/20
Stackdriver Logging	Enabled
Stackdriver Monitoring	Enabled
Private cluster	Disabled
Master authorized networks	Disabled
Network policy	Disabled
Legacy authorization	Disabled
Maintenance window	Any time
Cloud TPU	Disabled

Labels

None

[⌵ Add-ons](#)

[⌵ Permissions](#)

Node Pools

Node pools are separate instance groups running Kubernetes in a cluster. You may add node pools in different zones for higher availability, or add node pools of different type machines. To add a node pool, click Edit. [Learn more](#)

```

Cloud Shell
(gsw-k8s-3) (gsw-k8s-3) x +
Every 1.0s: gcloud container clusters list
cs-6000-devshell-vm-e3cee62-45c6-426f-892b-9bd785082044: Mon Sep  3 16:15:39 2018
NAME      LOCATION  MASTER_VERSION  MASTER_IP      MACHINE_TYPE  NODE_VERSION  NUM_NODES  STATUS
cluster-1 us-east1-b 1.10.6-gke.2    35.237.54.93   n1-standard-2 1.10.6-gke.2 4          STOPPING
cluster-2 us-east1-b 1.10.6-gke.2    35.237.47.212 n1-standard-2 1.10.6-gke.2 4          RUNNING

```

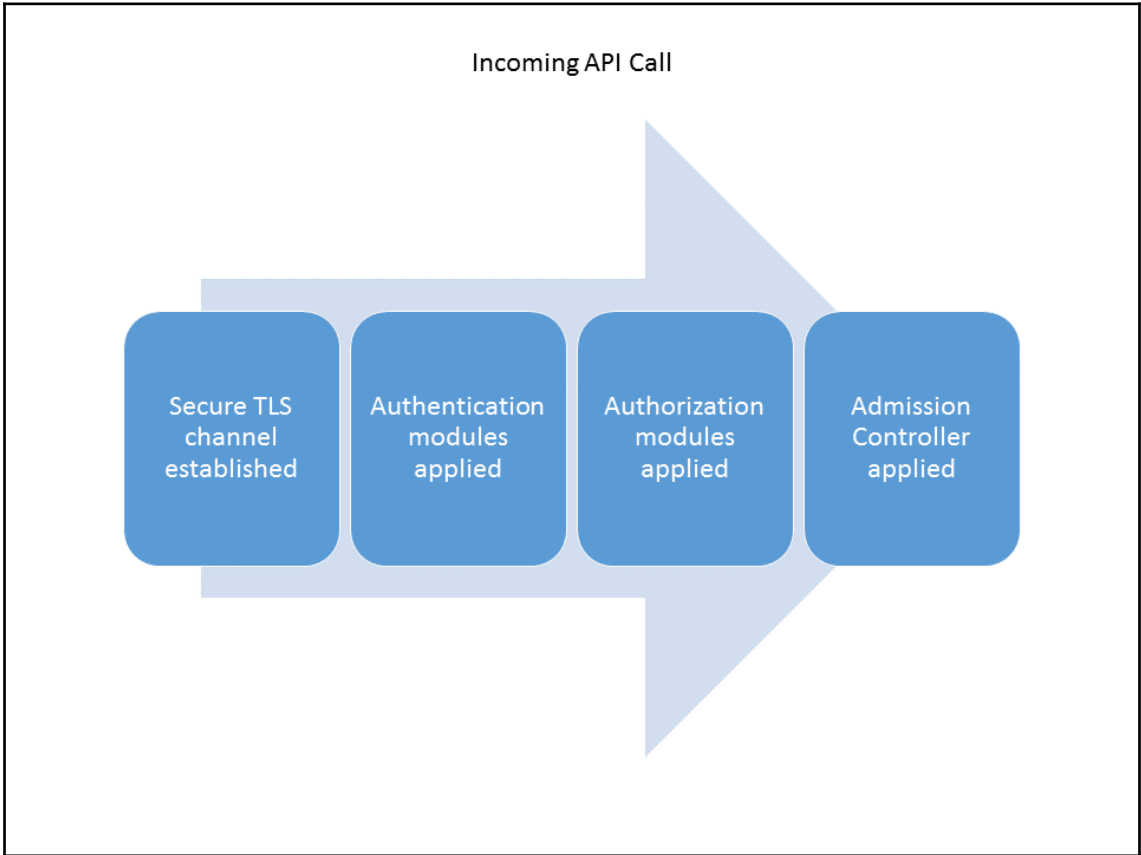
Google Cloud Platform GSW K8s 3

Compute Engine VM instances CREATE INSTANCE IMPORT VM REFRESH

Filter VM instances

<input type="checkbox"/>	Name ^	Zone	Recommendation	Internal IP	External IP	Connect
<input type="checkbox"/>	gke-cluster-2-default-pool-5db3e967-3q5h	us-east1-b		10.142.0.8 (nic0)	35.237.45.241	SSH
<input type="checkbox"/>	gke-cluster-2-default-pool-5db3e967-4crt	us-east1-b		10.142.0.7 (nic0)	35.229.107.14	SSH
<input type="checkbox"/>	gke-cluster-2-default-pool-5db3e967-4zqz	us-east1-b		10.142.0.6 (nic0)	35.229.84.39	SSH
<input type="checkbox"/>	gke-cluster-2-default-pool-5db3e967-zwvq	us-east1-b		10.142.0.9 (nic0)	35.227.91.178	SSH

Chapter 13: Cluster Authentication, Authorization, and Container Security



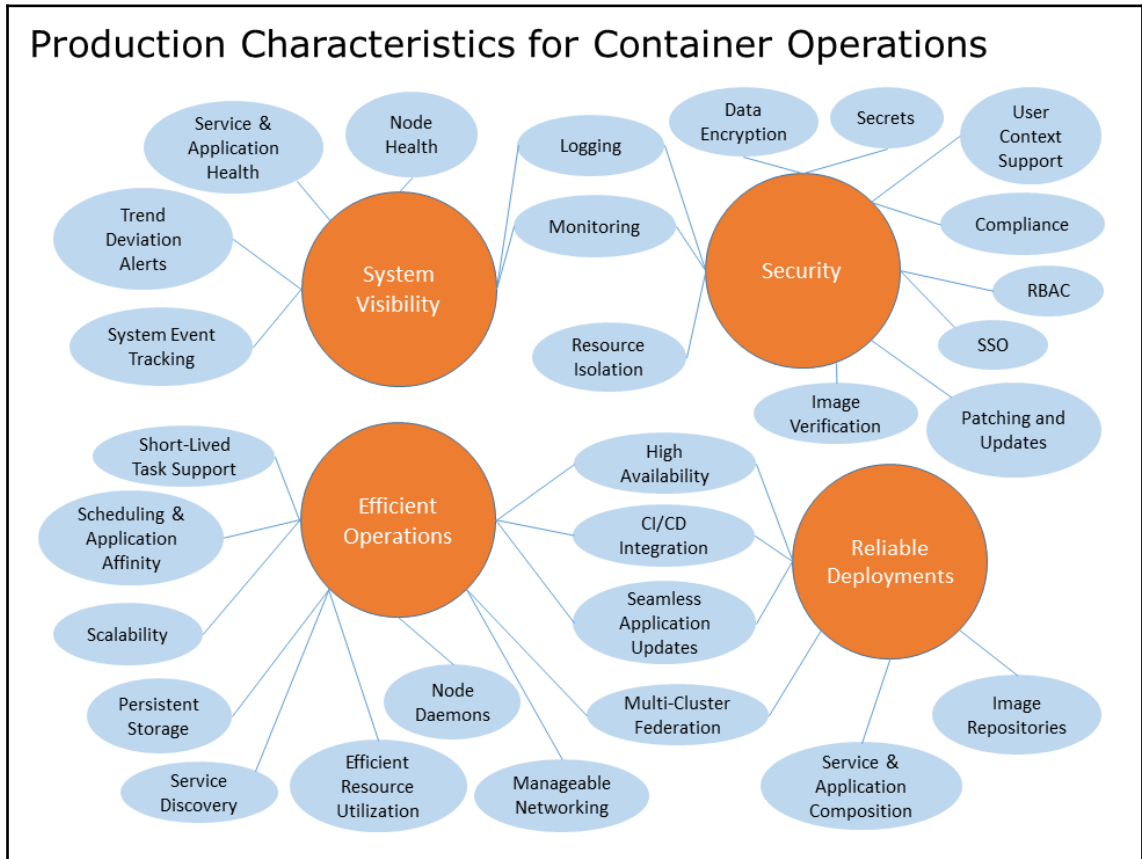
NAME	DATA	CAPS	SELINUX	RUNASUSER	FSGROUP
<code>gce.event-exporter</code>	<code>false</code>		<code>RunAsAny</code>	<code>RunAsAny</code>	<code>RunAsAny</code>
<code>gce.fluentd-gcp</code>	<code>false</code>		<code>RunAsAny</code>	<code>RunAsAny</code>	<code>RunAsAny</code>
<code>gce.persistent-volume-binder</code>	<code>false</code>		<code>RunAsAny</code>	<code>RunAsAny</code>	<code>RunAsAny</code>
<code>gce.privileged</code>	<code>true</code>	*	<code>RunAsAny</code>	<code>RunAsAny</code>	<code>RunAsAny</code>
<code>gce.unprivileged-addon</code>	<code>false</code>		<code>RunAsAny</code>	<code>RunAsAny</code>	<code>RunAsAny</code>

```

Name:          myroottest-588dfdcf85
Namespace:     default
Selector:      pod-template-hash=1448987941,run=myroottest
Labels:        pod-template-hash=1448987941
                run=myroottest
Annotations:   deployment.kubernetes.io/desired-replicas=1
                deployment.kubernetes.io/max-replicas=2
                deployment.kubernetes.io/revision=1
Controlled By: Deployment/myroottest
Replicas:      0 current / 1 desired
Pods Status:   0 Running / 0 Waiting / 0 Succeeded / 0 Failed
Pod Template:
  Labels:       pod-template-hash=1448987941
                run=myroottest
  Containers:
    myroottest:
      Image:      jonbaier/node-express-info:latest
      Port:       <none>
      Host Port:  <none>
      Environment: <none>
      Mounts:     <none>
      Volumes:    <none>
  Conditions:
    Type           Status  Reason
    ----           -
    ReplicaFailure True    FailedCreate
Events:
  Type           Reason          Age             From              Message
  ----           -
  Warning        FailedCreate    10s (x14 over 51s) replicaset-controller Error creating: pods "myroottest-588dfdcf85-"
is forbidden: unable to validate against any
pod security policy: []

```

Chapter 14: Hardening Kubernetes



Chapter 15: Kubernetes Infrastructure Management

IaaS Provider	Config. Mgmt.	OS	Networking
any	any	multi-support	any CNI
Google Kubernetes Engine			GCE
Stackpoint.io		multi-support	multi-support
AppsCode.com	Saltstack	Debian	multi-support
Madcore.Ai	Jenkins DSL	Ubuntu	flannel
Platform9		multi-support	multi-support
Kublr	custom	multi-support	multi-support
Kubernatic		multi-support	multi-support
IBM Cloud Kubernetes Service		Ubuntu	IBM Cloud Networking + Calico
Giant Swarm		CoreOS	flannel and/or Calico
GCE	Saltstack	Debian	GCE
Azure Kubernetes Service		Ubuntu	Azure
Azure (IaaS)		Ubuntu	Azure
Bare-metal	custom	Fedora	<i>none</i>
Bare-metal	custom	Fedora	flannel

