

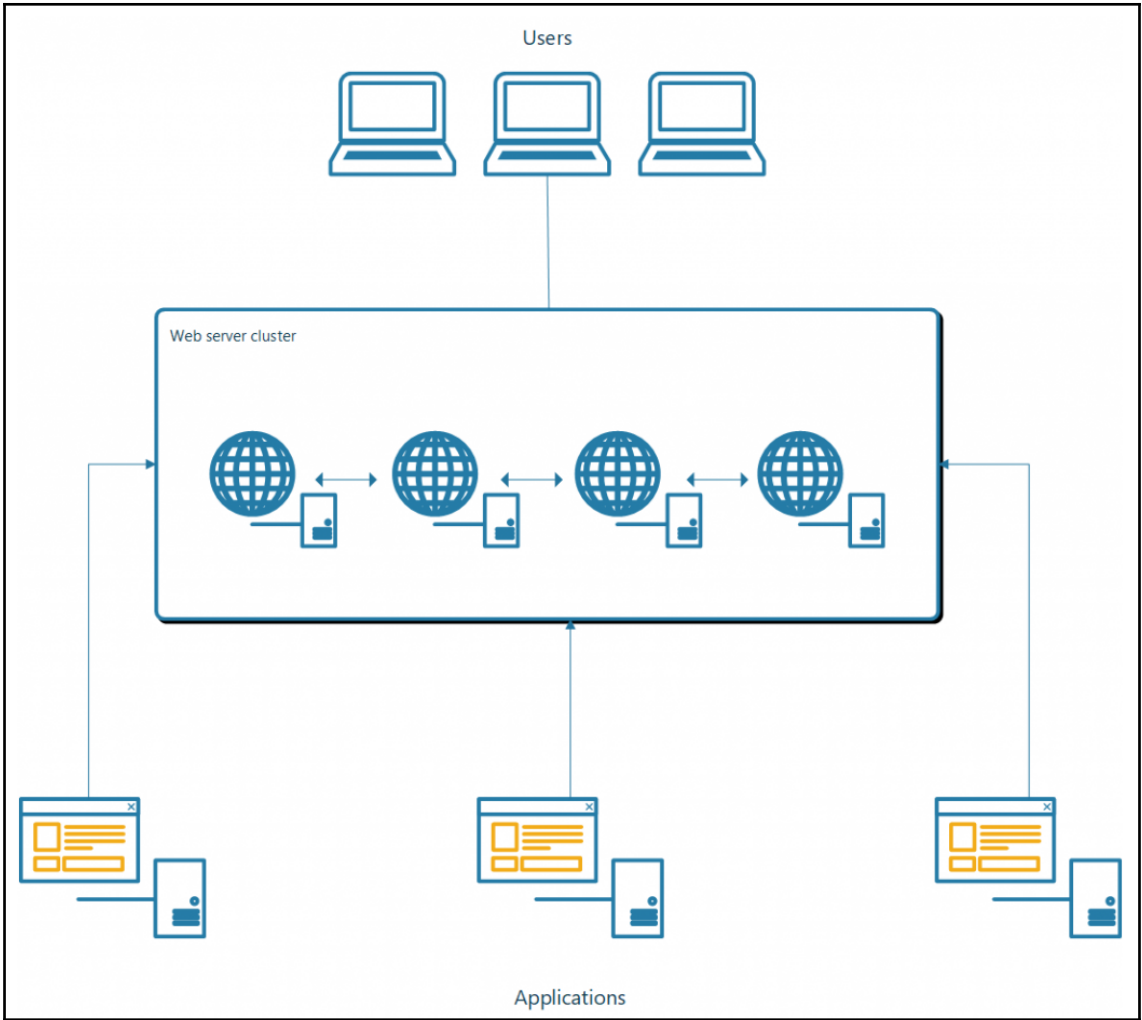


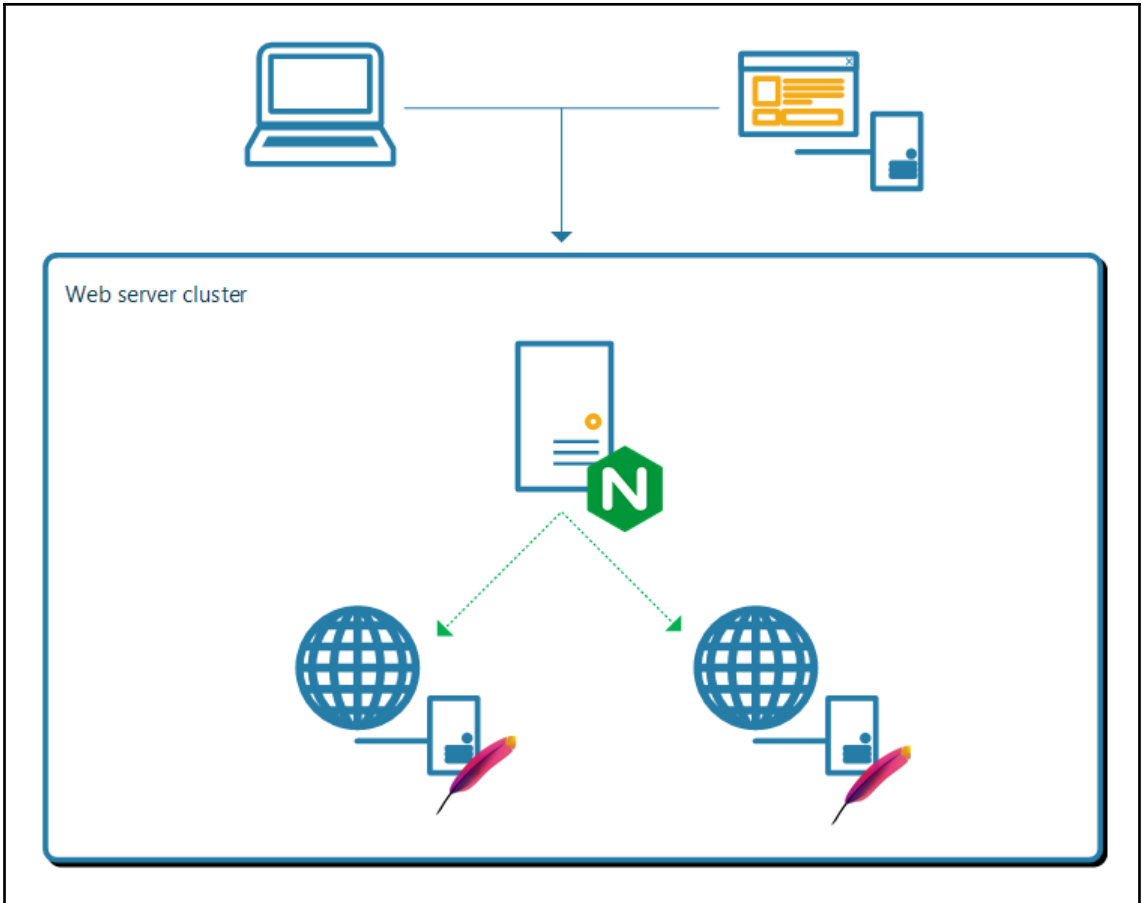
Table of Contents

	1
Index	74

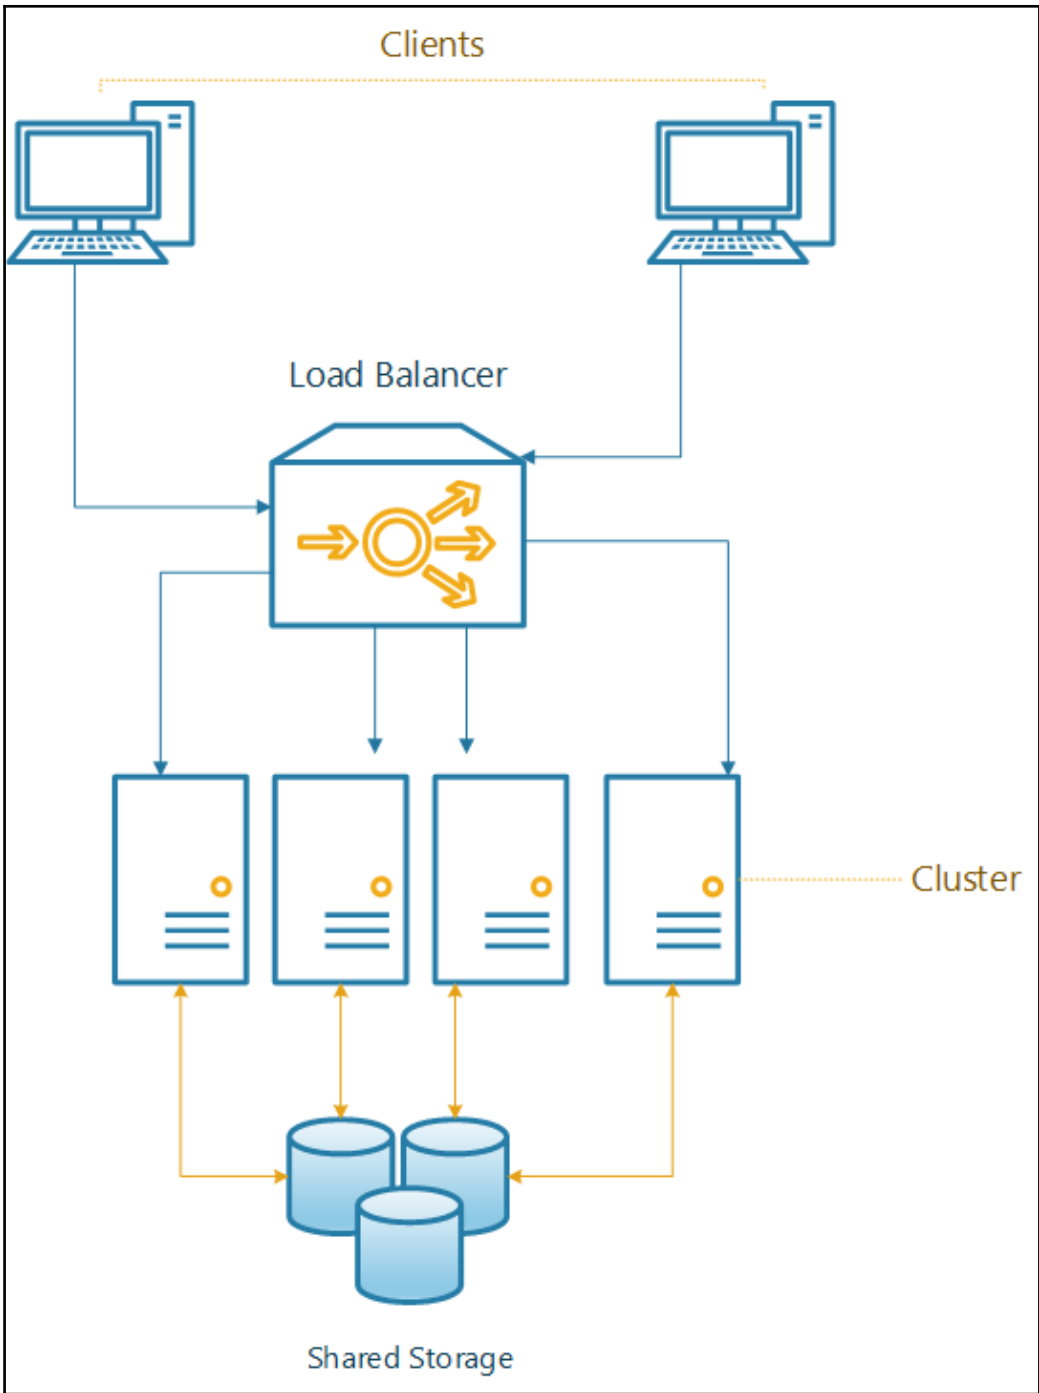
Chapter 1: Introduction to Design Methodology










Chapter 2: Defining GlusterFS Storage





 **174 Contributors**

 **11,758 Commits**




Block Storage

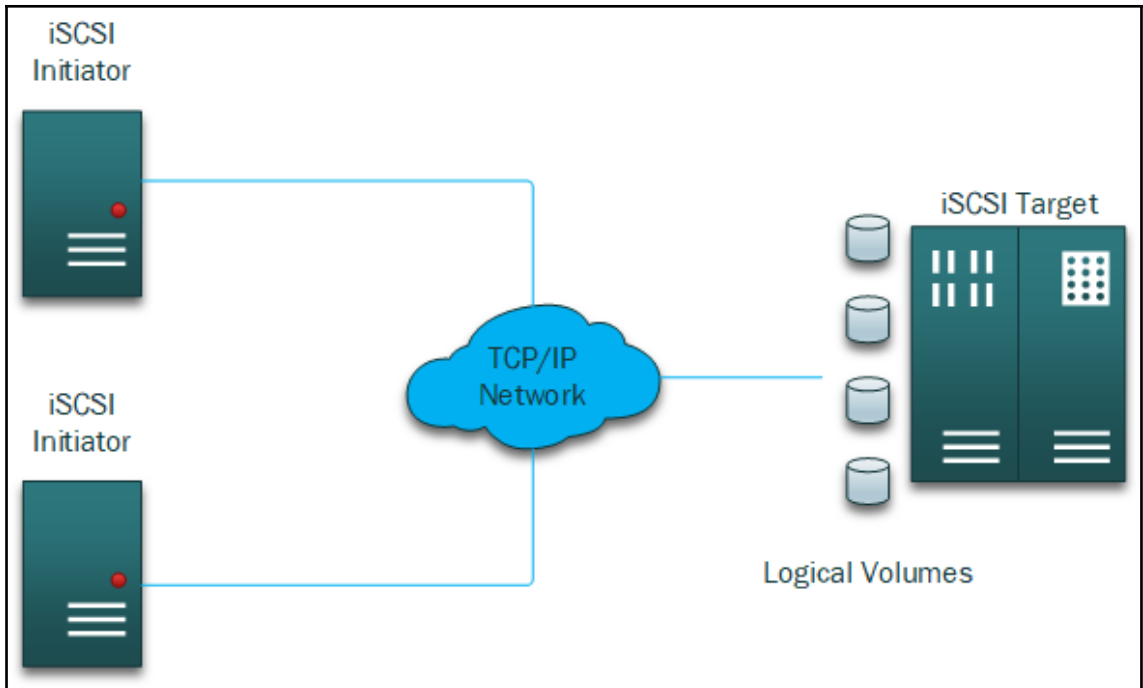


File Storage

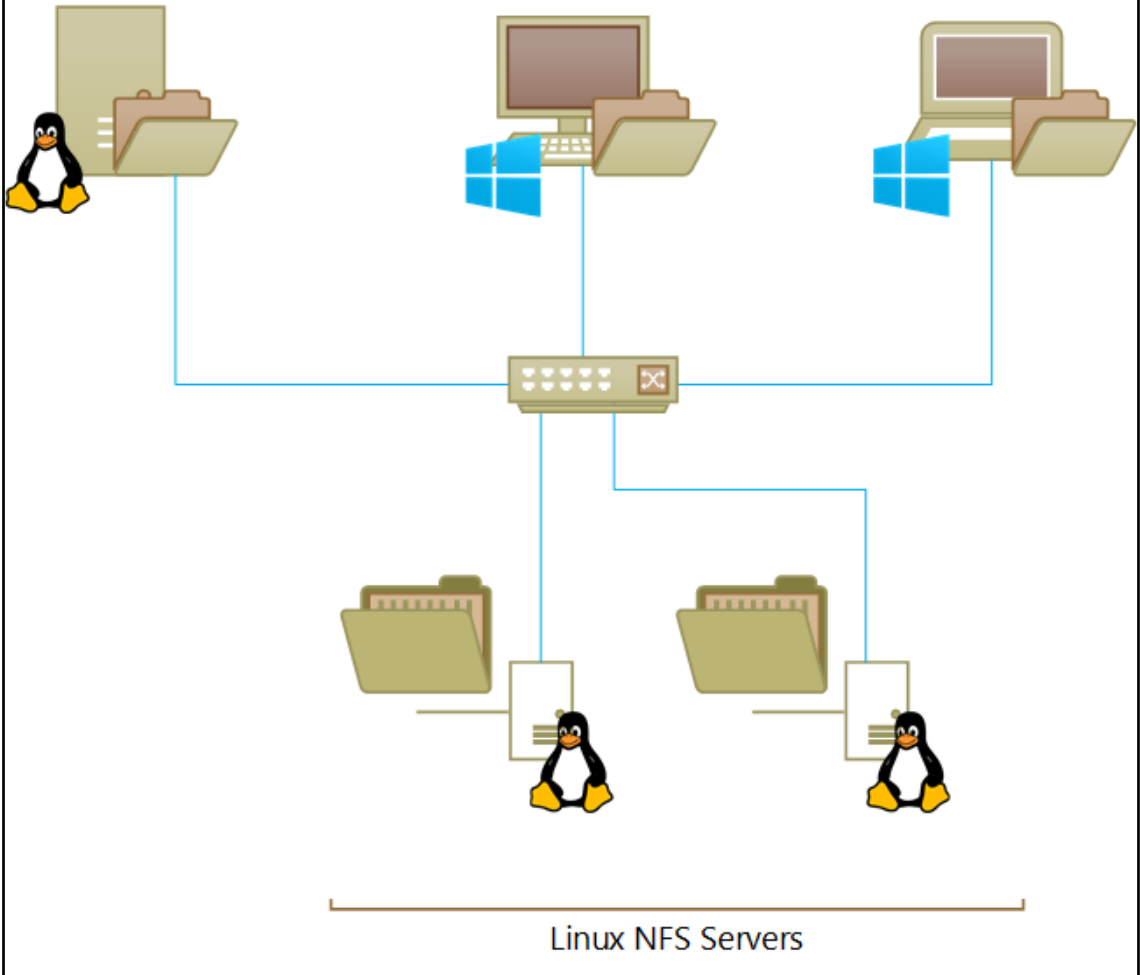


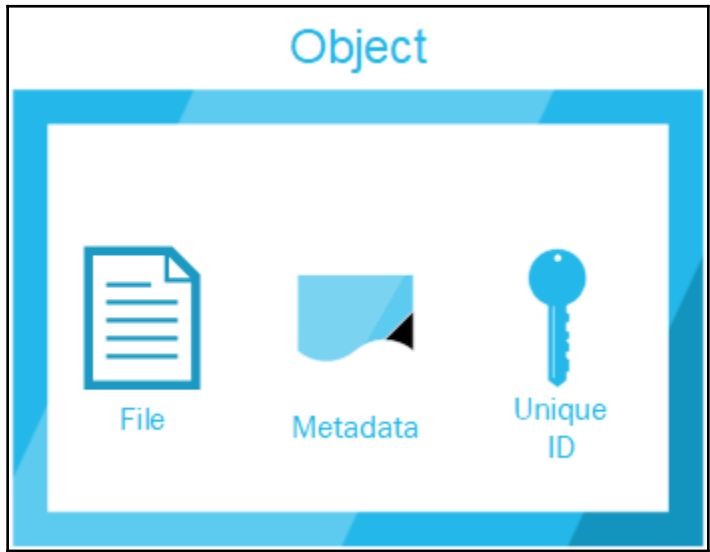
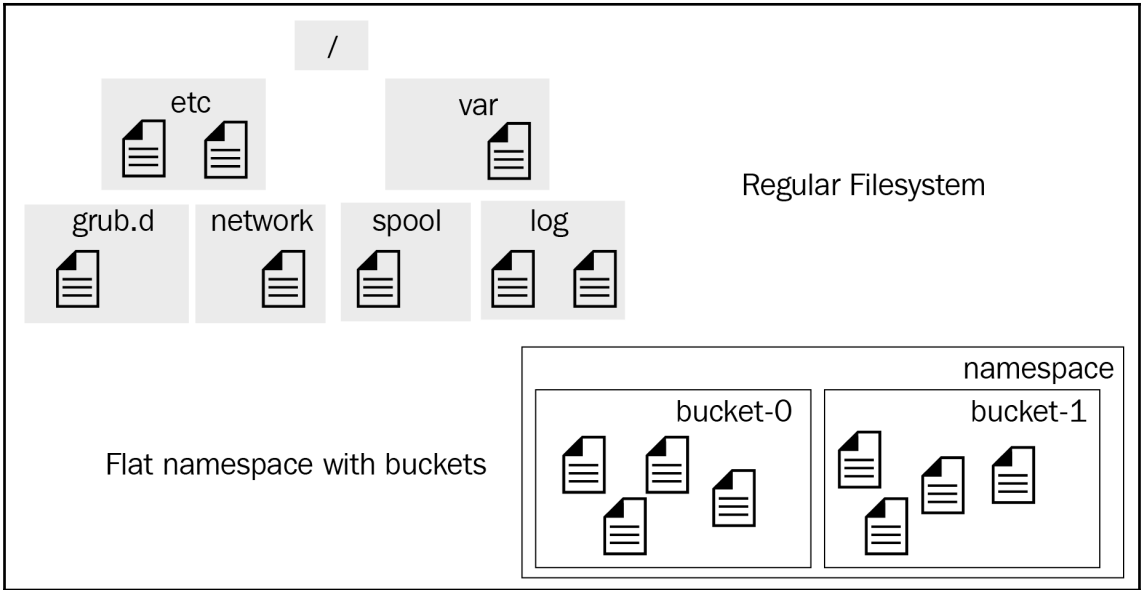
Object Storage


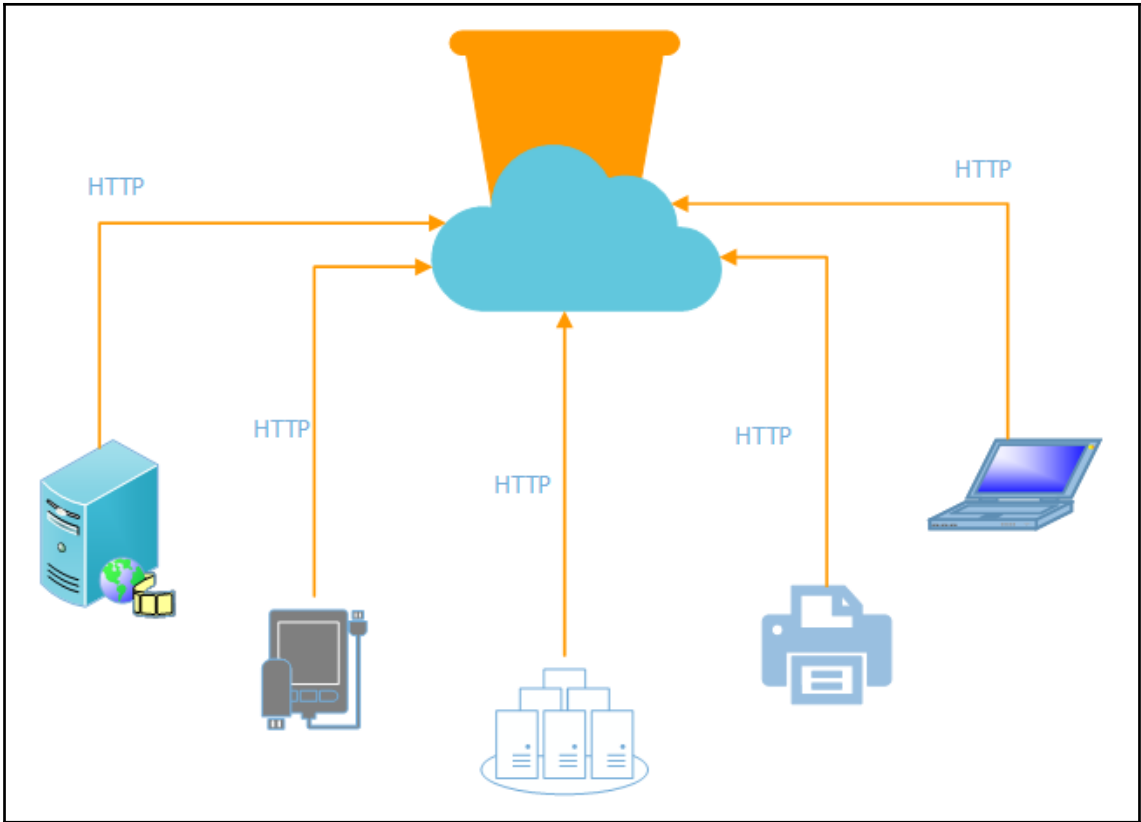


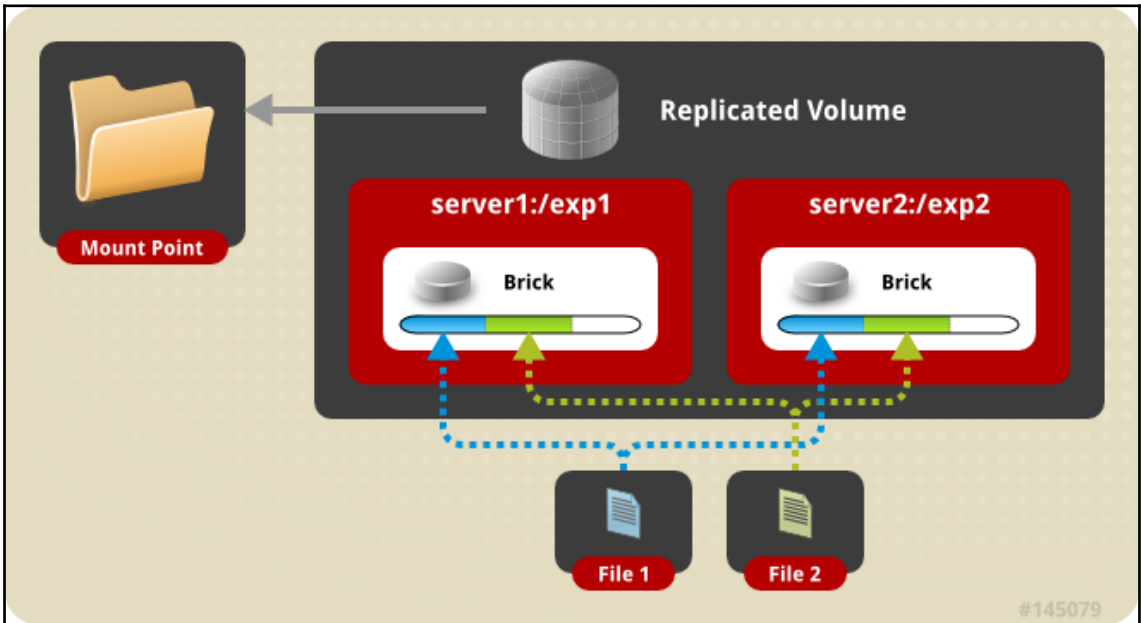
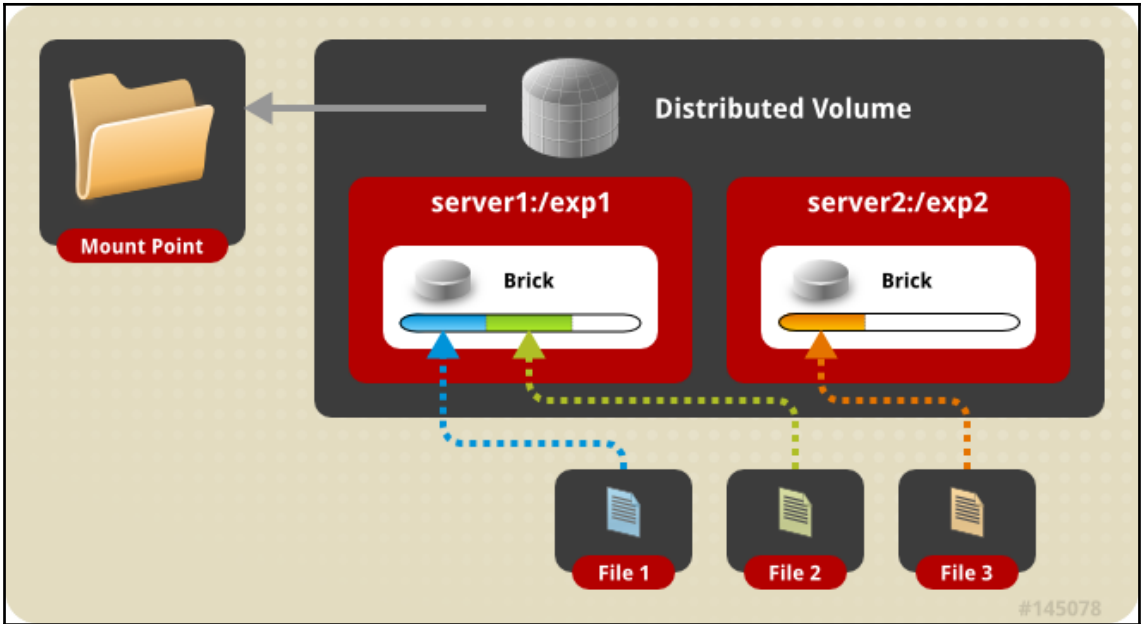


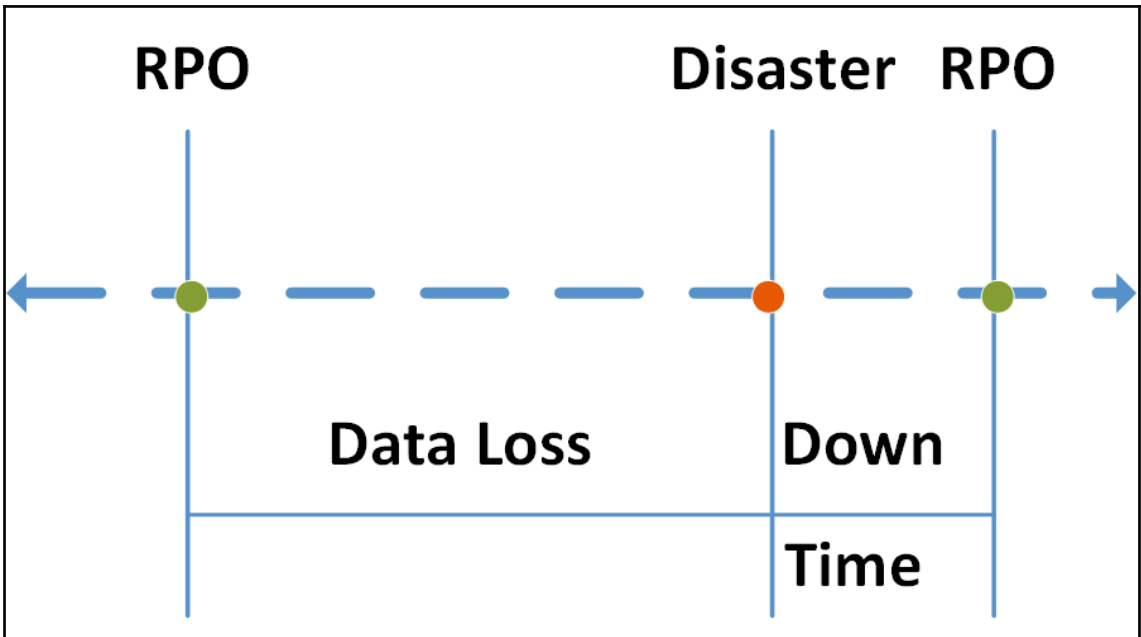
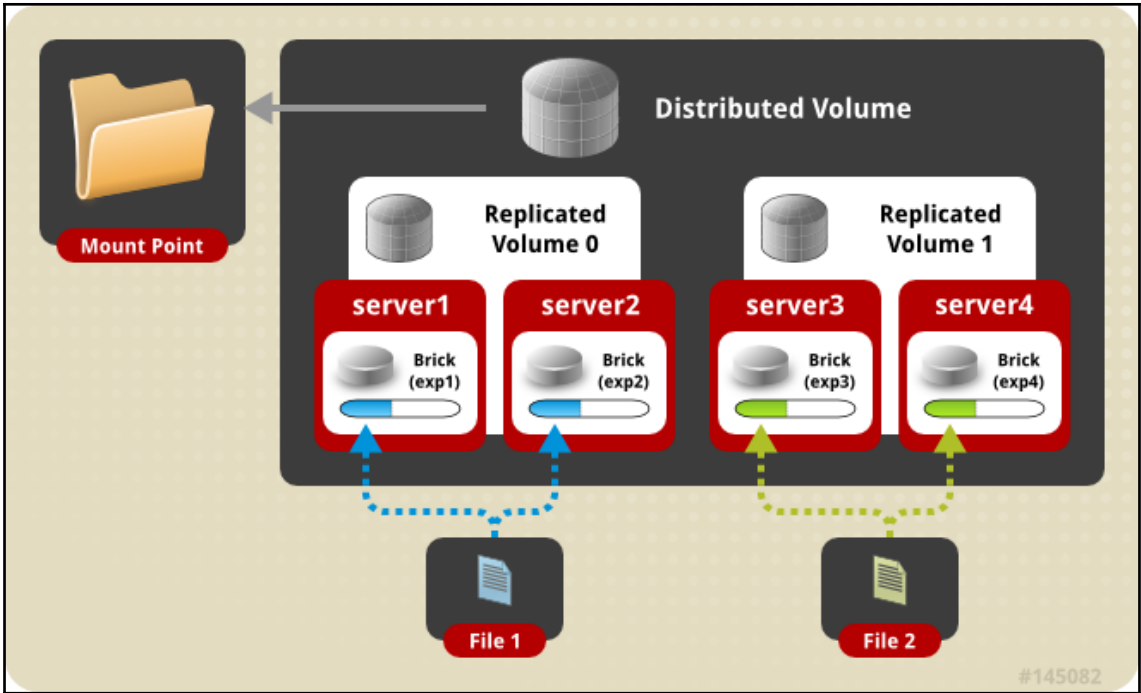
Different OS Clients Accessing the NFS Export



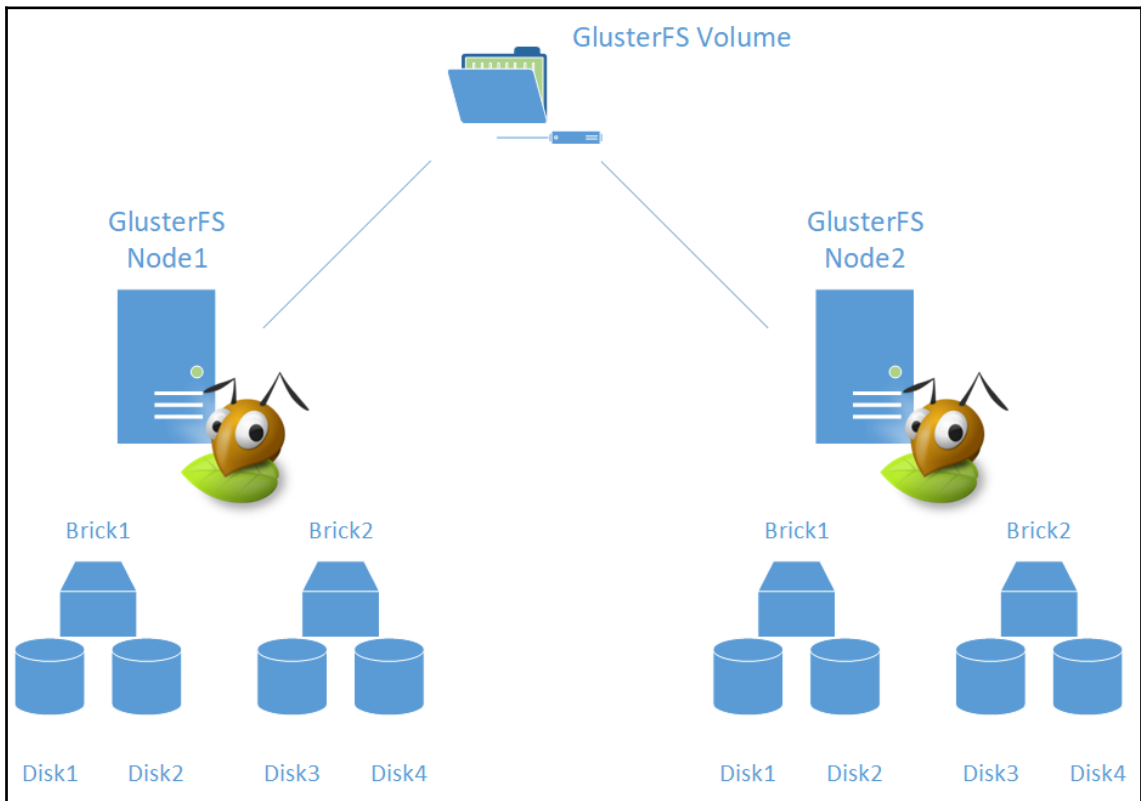


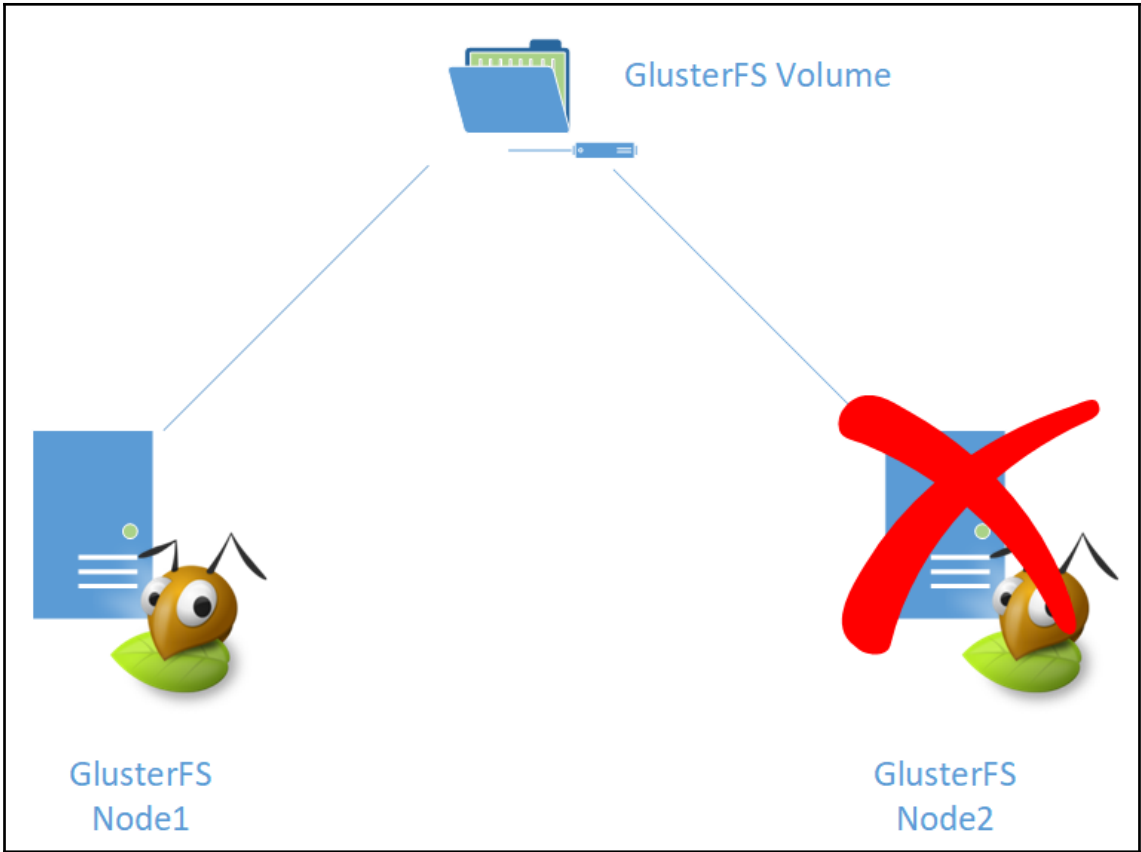


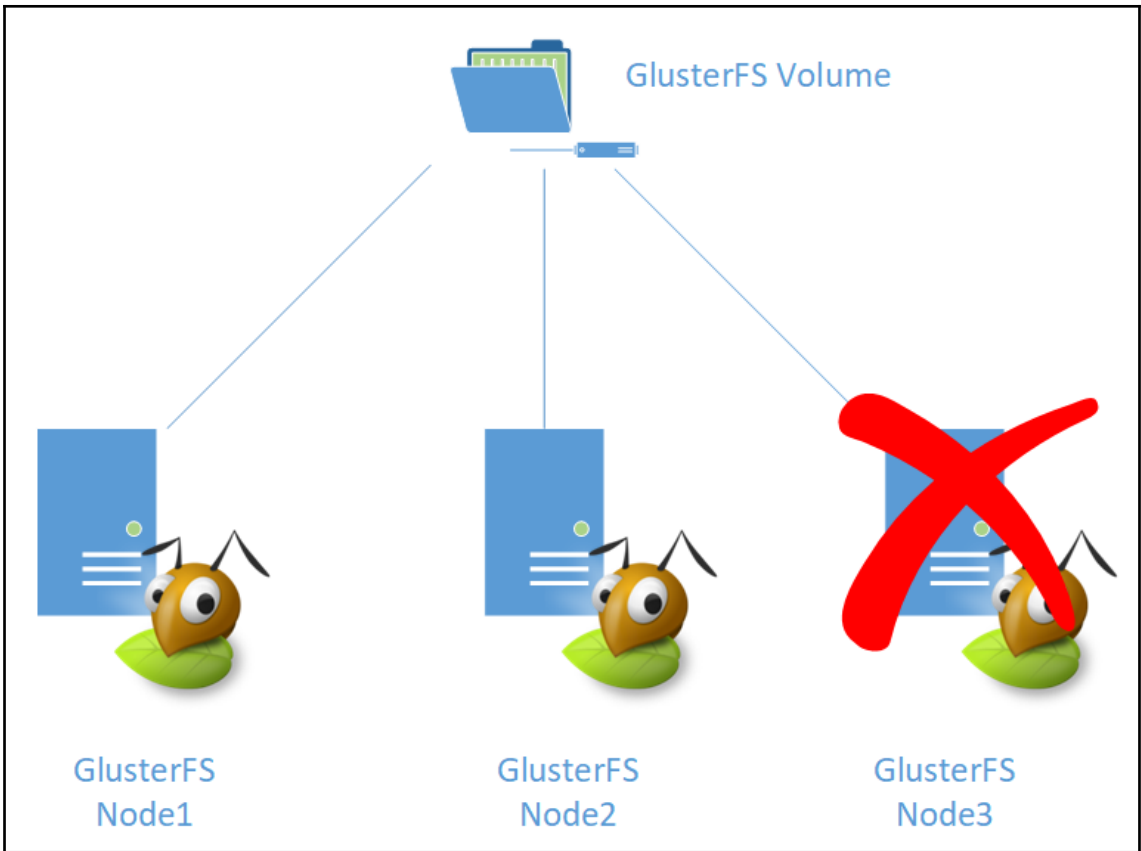


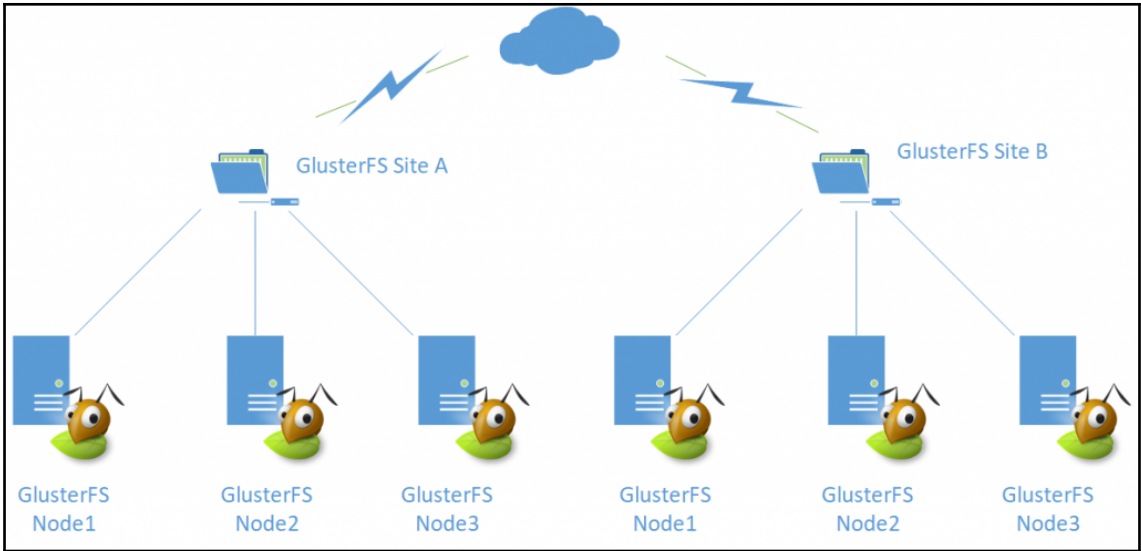


Chapter 3: Architecting a Storage Cluster







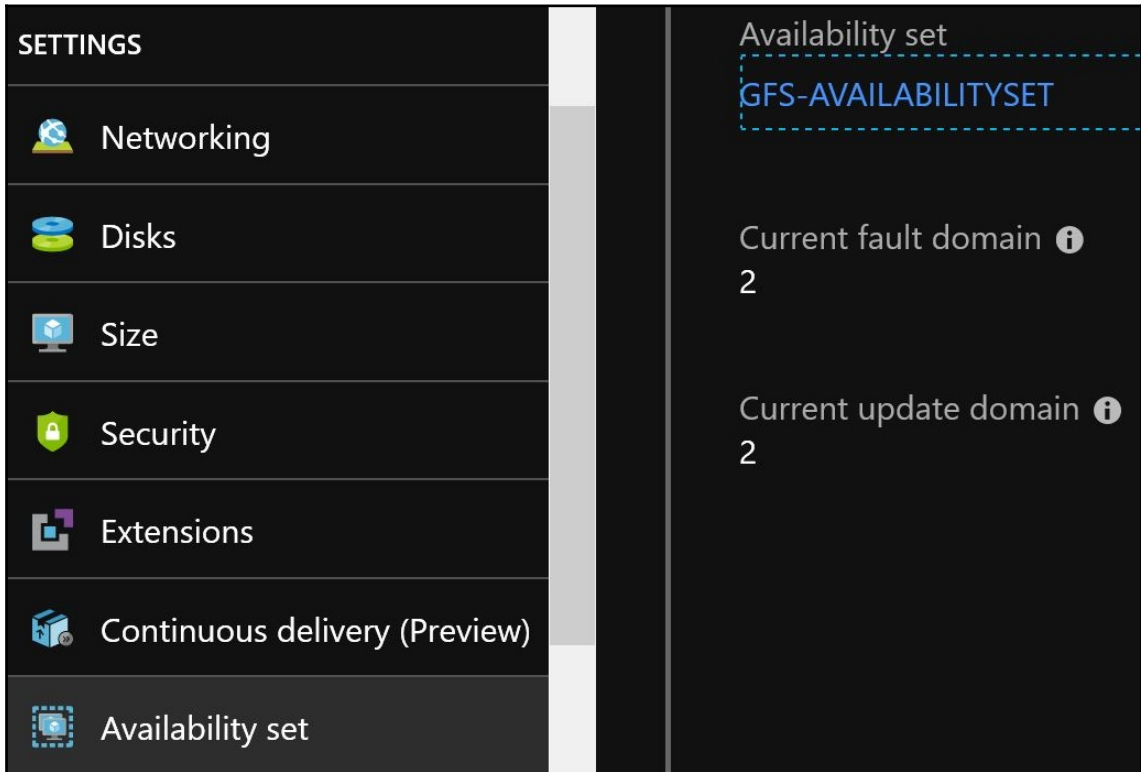


```

08/01/2018 12:49:19 PM
avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.22    0.00   1.54    0.29    0.00   97.94

Device            r/s     w/s   rMB/s   wMB/s   rrqm/s   wrqm/s   %rrqm   %wrqm  r_await  w_await  aqu-sz  rareq-sz  wareq-sz  svctm   %util
sda                1.42   20.47    0.03    0.22    0.03    0.15    2.08    0.71   24.60    0.87    0.05   24.57    11.08    1.40    3.06
sdb                3.54    7.17    0.09    0.51    0.04    2.16    1.22   23.15    5.37    8.47    0.08   24.95    72.30    0.50    0.53
sdc                1.42   20.43    0.03    0.22    0.03    0.15    2.40    0.72   24.56    0.88    0.05   24.60    11.11    1.40    3.07
sdd                1.42   20.51    0.03    0.22    0.03    0.14    2.33    0.69   24.45    0.86    0.05   24.57    11.06    1.39    3.06
sde                1.43   19.81    0.03    0.22    0.04    0.15    2.53    0.75   24.42    1.12    0.06   24.56    11.46    1.47    3.13
sdf                1.43   19.76    0.03    0.22    0.03    0.15    2.35    0.75   24.74    1.12    0.06   24.66    11.49    1.49    3.15
loop0              0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.25    0.00    0.00   25.05    0.00    0.17    0.00
loop1              0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.21    1.83    0.00   22.64    1.33    0.08    0.00
  
```

Chapter 4: Using GlusterFS on the Cloud Infrastructure



Disk	Size GB	Price/month	Price/GB/month
S4	32	\$1.54	\$0.05
S6	64	\$3.01	\$0.05
S10	128	\$5.89	\$0.05
S15	256	\$11.33	\$0.04
S20	512	\$21.76	\$0.04
S30	1024	\$40.96	\$0.04
S40	2048	\$81.92	\$0.04
S50	4096	\$163.84	\$0.04

```
[root@gfs1 /]# zpool status
pool: brick1
state: ONLINE
scan: none requested
config:
```

NAME	STATE	READ	WRITE	CKSUM
brick1	ONLINE	0	0	0
scsi-360022480b73cfc2c8ec934cab5ddc3cf	ONLINE	0	0	0
scsi-36002248035a747697e238a3c1143b742	ONLINE	0	0	0
scsi-360022480f0da979b536cde32a4a17406	ONLINE	0	0	0
scsi-360022480fb9d18bbdfb9175fd3e0bbf2	ONLINE	0	0	0

```
errors: No known data errors
```

```
[root@gfs1 /]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda2       30G   1.8G   28G   6% /
devtmpfs        16G    0    16G   0% /dev
tmpfs           16G    0    16G   0% /dev/shm
tmpfs           16G   9.1M   16G   1% /run
tmpfs           16G    0    16G   0% /sys/fs/cgroup
/dev/sda1       497M  105M  393M  22% /boot
tmpfs           3.2G    0   3.2G   0% /run/user/1000
brick1          493G   33G  461G   7% /bricks/brick1
[root@gfs1 /]#
```

```
[root@gfs1 /]# zpool status
```

```
pool: brick1
state: ONLINE
scan: none requested
config:
```

NAME	STATE	READ	WRITE	CKSUM
brick1	ONLINE	0	0	0
scsi-360022480b73cfc2c8ec934cab5ddc3cf	ONLINE	0	0	0
scsi-36002248035a747697e238a3c1143b742	ONLINE	0	0	0
scsi-360022480f0da979b536cde32a4a17406	ONLINE	0	0	0
scsi-360022480fb9d18bbdfb9175fd3e0bbf2	ONLINE	0	0	0
cache				
sdb	ONLINE	0	0	0

```
errors: No known data errors
```

```
[root@gfs1 /]#
```

```
[root@gfs1 /]# systemctl enable glusterd
[root@gfs1 /]# systemctl start glusterd.service
[root@gfs1 /]# systemctl status glusterd
● glusterd.service - GlusterFS, a clustered file-system server
   Loaded: loaded (/usr/lib/systemd/system/glusterd.service; enabled; vendor preset: disabled)
   Active: active (running) since Wed 2018-08-15 04:09:14 UTC; 4s ago
     Process: 4276 ExecStart=/usr/sbin/glusterd -p /var/run/glusterd.pid --log-level $LOG_LEVEL $GLUSTERD_OPTIONS
    de=exited, status=0/SUCCESS
   Main PID: 4277 (glusterd)
      CGroup: /system.slice/glusterd.service
             └─4277 /usr/sbin/glusterd -p /var/run/glusterd.pid --log-level INFO

Aug 15 04:09:14 gfs1 systemd[1]: Starting GlusterFS, a clustered file-system server...
Aug 15 04:09:14 gfs1 systemd[1]: Started GlusterFS, a clustered file-system server.
[root@gfs1 /]#
```

```
[root@gfs1 /]# gluster peer probe gfs2
peer probe: success.
[root@gfs1 /]# gluster peer probe gfs3
peer probe: success.
[root@gfs1 /]# gluster peer status
Number of Peers: 2

Hostname: gfs2
Uuid: 2d0ee0b9-3a6f-449e-8945-64f1aa35b962
State: Peer in Cluster (Connected)

Hostname: gfs3
Uuid: 80a06eae-fb2a-4a65-a7d4-91ed9a91c050
State: Peer in Cluster (Connected)
[root@gfs1 /]#
```

```

[root@gfs1 /]# gluster volume create gvol1 disperse 3 gfs{1..3}:/bricks/brick1/gvol1
volume create: gvol1: success: please start the volume to access data
[root@gfs1 /]# gluster volume start gvol1
volume start: gvol1: success
[root@gfs1 /]# gluster volume status gvol1
Status of volume: gvol1
Gluster process                                TCP Port  RDMA Port  Online  Pid
-----
Brick gfs1:/bricks/brick1/gvol1                49152     0           Y       7541
Brick gfs2:/bricks/brick1/gvol1                49152     0           Y       5694
Brick gfs3:/bricks/brick1/gvol1                49152     0           Y       5145
Self-heal Daemon on localhost                  N/A       N/A         Y       7625
Self-heal Daemon on gfs2                      N/A       N/A         Y       5768
Self-heal Daemon on gfs3                      N/A       N/A         Y       5216

Task Status of Volume gvol1
-----
There are no active volume tasks

[root@gfs1 /]#

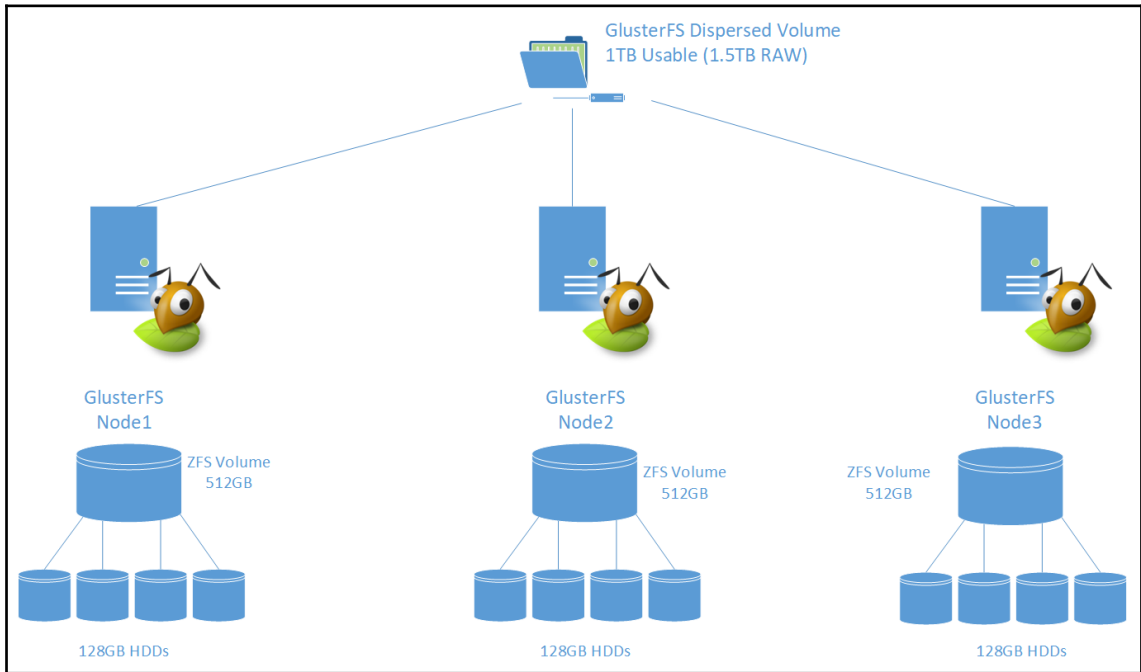
```

```

options zfs zfs_arc_max=27917287424
options zfs l2arc_write_max=2621440000

```

Chapter 5: Analyzing Performance in a Gluster System




```
[root@gfs1 /]# zpool iostat -vLP 1
```

pool	capacity		operations		bandwidth	
	alloc	free	read	write	read	write
brick1	62.3G	446G	0	1	11.3K	26.0K
/dev/sdf1	15.4G	112G	0	0	2.61K	6.18K
/dev/sdd1	15.9G	111G	0	0	2.33K	6.73K
/dev/sdc1	15.6G	111G	0	0	2.70K	6.20K
/dev/sde1	15.4G	112G	0	0	3.65K	6.92K
cache	-	-	-	-	-	-
/dev/sdb1	8K	678G	0	0	303	98

```
08/29/2018 01:16:57 AM
avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           0.17    0.00    1.61    0.22    0.00    98.00

Device      r/s    w/s   rMB/s   wMB/s  rrqm/s  wrqm/s  %rrqm  %wrqm  r_await  w_await
sda         0.77   19.07   0.01    0.14   0.00    0.01   0.61   0.08   37.69   0.49
sdb         2.30    5.49   0.01    0.31   0.01    0.47   0.46   7.86    0.59   0.50
sdc         0.78   19.06   0.01    0.14   0.00    0.01   0.55   0.08   37.61   0.49
sdd         0.78   19.05   0.01    0.14   0.00    0.01   0.56   0.07   37.43   0.49
sde         0.77   18.59   0.01    0.14   0.00    0.02   0.59   0.09   39.46   0.70
sdf         0.78   18.53   0.01    0.14   0.00    0.02   0.60   0.08   38.07   0.70
```

```
fio-3.1
Starting 16 processes
jobs: 16 (f=16): [m(16)][12.2%][r=134MiB/s,w=145MiB/s][r=134,w=145 IOPS][eta 02m:38s]
```

```

file1: (groupid=0, jobs=16): err= 0: pid=2747: Wed Aug 29 05:18:42 2018
read: IOPS=141, BW=141MiB/s (148MB/s)(8482MiB/60155msec)
  slat (msec): min=7, max=387, avg=91.82, stdev=55.91
  clat (usec): min=7, max=4742.4k, avg=3422222.67, stdev=595136.99
  lat (msec): min=44, max=4790, avg=3514.04, stdev=596.21
  clat percentiles (msec):
  | 1.00th=[ 584], 5.00th=[ 2702], 10.00th=[ 2937], 20.00th=[ 3138],
  | 30.00th=[ 3272], 40.00th=[ 3373], 50.00th=[ 3473], 60.00th=[ 3574],
  | 70.00th=[ 3708], 80.00th=[ 3842], 90.00th=[ 4010], 95.00th=[ 4144],
  | 99.00th=[ 4396], 99.50th=[ 4463], 99.90th=[ 4597], 99.95th=[ 4665],
  | 99.99th=[ 4732]
  bw ( KiB/s): min= 2043, max=22894, per=6.32%, avg=9125.00, stdev=3924.97, samples=1799
  iops       : min=   1, max=  22, avg=  8.83, stdev=  3.82, samples=1799
write: IOPS=143, BW=144MiB/s (151MB/s)(8649MiB/60155msec)
  slat (usec): min=1515, max=211350, avg=20973.18, stdev=21320.17
  clat (usec): min=9, max=4742.2k, avg=3442255.32, stdev=599417.47
  lat (msec): min=28, max=4760, avg=3463.23, stdev=599.80
  clat percentiles (msec):
  | 1.00th=[ 676], 5.00th=[ 2702], 10.00th=[ 2970], 20.00th=[ 3171],
  | 30.00th=[ 3272], 40.00th=[ 3406], 50.00th=[ 3507], 60.00th=[ 3608],
  | 70.00th=[ 3708], 80.00th=[ 3842], 90.00th=[ 4044], 95.00th=[ 4178],
  | 99.00th=[ 4396], 99.50th=[ 4463], 99.90th=[ 4665], 99.95th=[ 4665],
  | 99.99th=[ 4732]
  bw ( KiB/s): min= 1865, max=25026, per=6.31%, avg=9295.26, stdev=3981.09, samples=1798
  iops       : min=   1, max=  24, avg=  9.00, stdev=  3.87, samples=1798
  lat (usec) : 10=0.06%, 20=0.02%, 50=0.01%
  lat (msec) : 50=0.03%, 100=0.06%, 250=0.23%, 500=0.39%, 750=0.39%
  lat (msec) : 1000=0.40%, 2000=1.55%, >=2000=96.86%
  cpu        : usr=0.08%, sys=1.46%, ctx=185112, majf=0, minf=470
  IO depths  : 1=0.1%, 2=0.2%, 4=0.4%, 8=0.7%, 16=1.5%, 32=3.0%, >=64=94.1%
  submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
  complete  : 0=0.0%, 4=99.9%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.1%, >=64=0.0%
  issued rw: total=8482,8649,0, short=0,0,0, dropped=0,0,0
  latency   : target=0, window=0, percentile=100.00%, depth=64

Run status group 0 (all jobs):
  READ: bw=141MiB/s (148MB/s), 141MiB/s-141MiB/s (148MB/s-148MB/s), io=8482MiB (8894MB), run=60155-60155msec
  WRITE: bw=144MiB/s (151MB/s), 144MiB/s-144MiB/s (151MB/s-151MB/s), io=8649MiB (9069MB), run=60155-60155msec

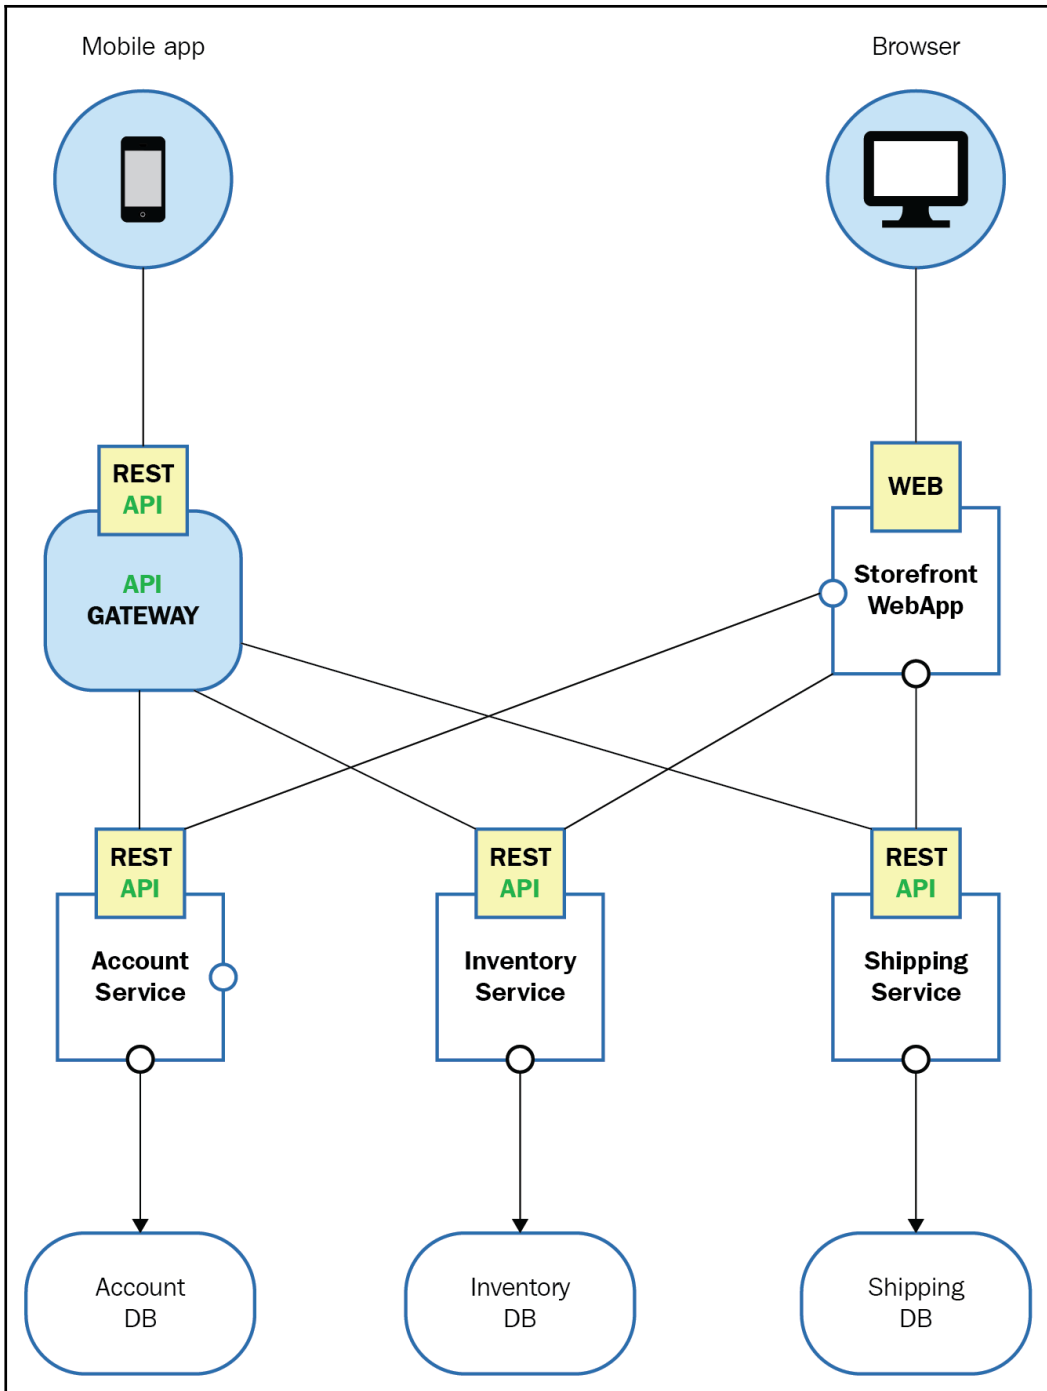
```

Gluster process	TCP Port	RDMA Port	Online	Pid
Brick gfs1:/bricks/brick1/gv011	49152	0	Y	2143
Brick gfs3:/bricks/brick1/gv011	49152	0	Y	2218
Self-heal Daemon on localhost	N/A	N/A	Y	1969
Self-heal Daemon on gfs3	N/A	N/A	Y	2080

Task Status of Volume gv011	Jobs
Jobs: 16 (f=16): [m(16)][62.4%][r=159MiB/s,w=169MiB/s][r=159,w=169 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][62.8%][r=135MiB/s,w=145MiB/s][r=135,w=145 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][63.2%][r=163MiB/s,w=177MiB/s][r=163,w=177 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][63.6%][r=161MiB/s,w=168MiB/s][r=161,w=168 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][64.0%][r=156MiB/s,w=149MiB/s][r=156,w=149 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][64.4%][r=172MiB/s,w=149MiB/s][r=172,w=149 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][64.8%][r=165MiB/s,w=141MiB/s][r=165,w=141 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][65.2%][r=159MiB/s,w=158MiB/s][r=159,w=158 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][65.6%][r=168MiB/s,w=152MiB/s][r=168,w=152 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][66.0%][r=176MiB/s,w=155MiB/s][r=176,w=155 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][66.4%][r=156MiB/s,w=171MiB/s][r=156,w=171 IOPS][eta 01m]	
Jobs: 16 (f=16): [m(16)][66.8%][r=154MiB/s,w=157MiB/s][r=154,w=157 IOPS][eta 01m]	

There are no active volume tasks

Chapter 6: Creating a Highly Available Self-Healing Architecture



	No ENTRYPOINT	ENTRYPOINT exec_entry p1_entry	ENTRYPOINT ["exec_entry", "p1_entry"]
No CMD	<i>error, not allowed</i>	/bin/sh -c exec_entry p1_entry	exec_entry p1_entry
CMD ["exec_cmd", "p1_cmd"]	exec_cmd p1_cmd	/bin/sh -c exec_entry p1_entry	exec_entry p1_entry exec_cmd p1_cmd
CMD ["p1_cmd", "p2_cmd"]	p1_cmd p2_cmd	/bin/sh -c exec_entry p1_entry	exec_entry p1_entry p1_cmd p2_cmd
CMD exec_cmd p1_cmd	/bin/sh -c exec_cmd p1_cmd	/bin/sh -c exec_entry p1_entry	exec_entry p1_entry /bin/sh -c exec_cmd p1_cmd

Scan results for **ubuntu:latest**

7 of 69 components are vulnerable Provide Feedback
Scanned a day ago

Layers

1 **ADD file:4bb62bb05874...7ab4c1fdece957a in /**
Compressed size: 30.2MB

COMPONENT	VULNERABILITY	SEVERITY
glibc 2.27-3ubuntu1 LGPL:lgpl License	CVE-2017-18269	Critical
	CVE-2018-11236	Major
	CVE-2018-11237	Major
systemd 237-3ubuntu10.3 LGPL:lgpl License	CVE-2018-6954	Critical
util-linux 2.31.1-0.4ubuntu3.1 GPL:copyleft License	CVE-2018-7738	Critical
berkeleydb 5.3.28-13.1ubuntu1 sleepycat:copyleft License	CVE-2016-0689	Major
	CVE-2016-0692	Major
	CVE-2016-0694	Major
	CVE-2016-3418	Major
	CVE-2016-0682	Major
shadow 4.5-1ubuntu1 BSD:permissive License	CVE-2018-7169	Major
selinux 2.7-2build2 Public domain:permissive License	CVE-2018-1063	Minor
coreutils 8.28-1ubuntu1 GPLv3:copyleft License	CVE-2017-18018	Minor

Components

Scan results for **alpine:latest**



Your image is clean! No known vulnerabilities were found.

[Provide Feedback](#)


Scanned 20 hours ago

Layers

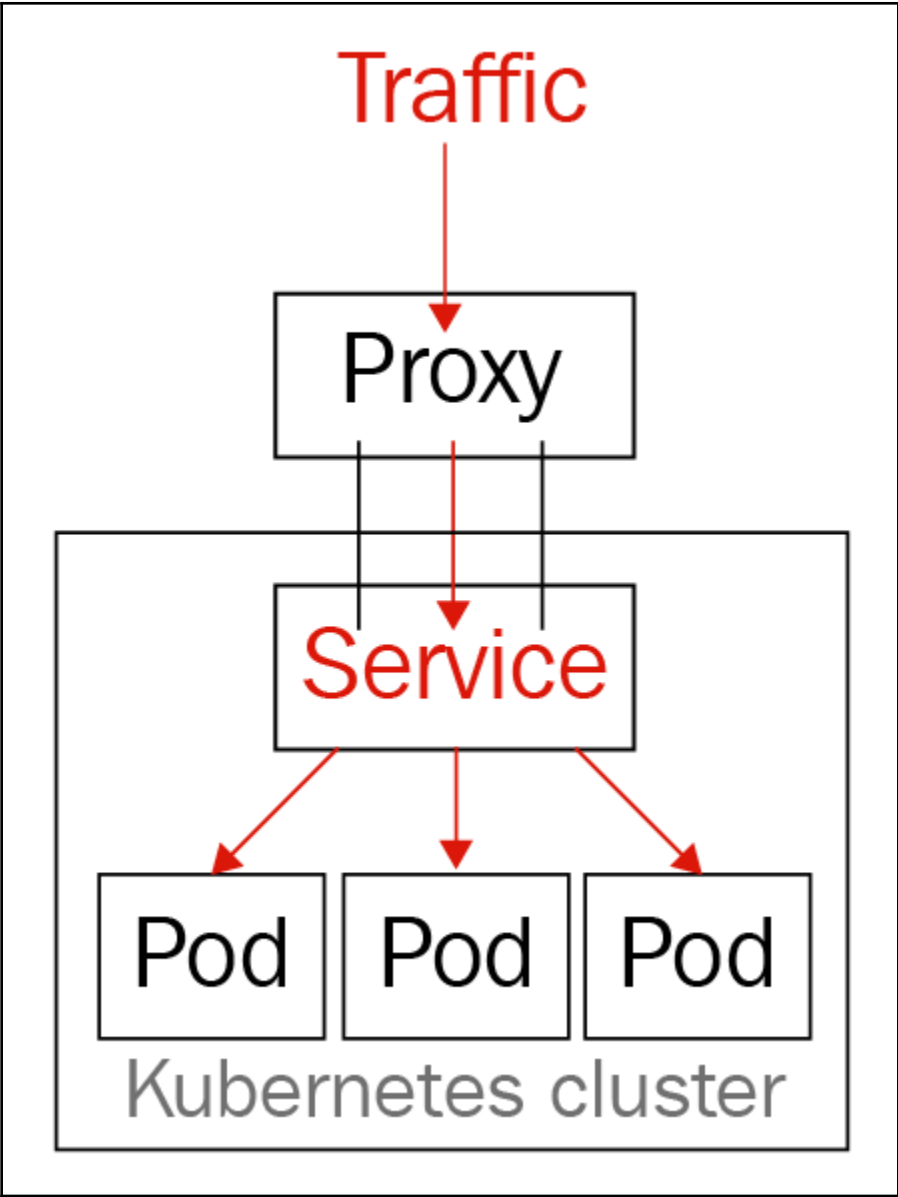
Components

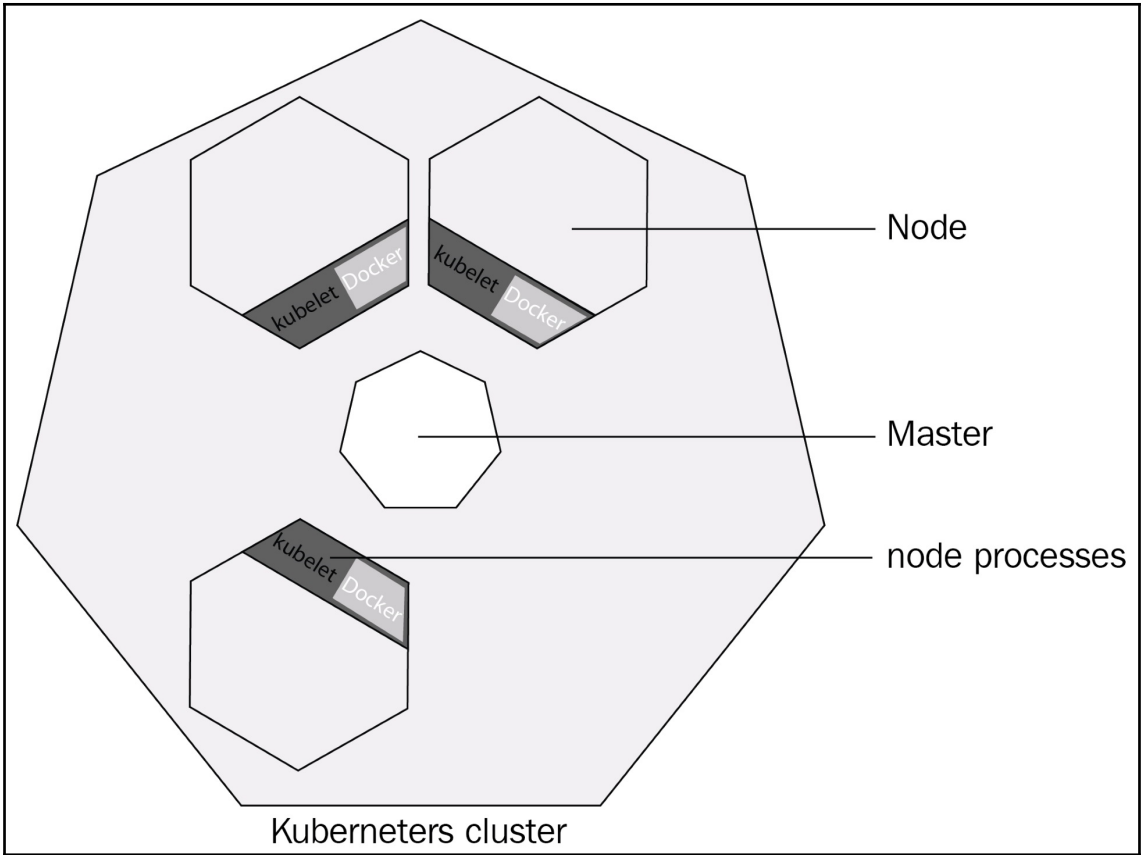
1 [ADD file:25f61d70254b...f6a33810393e95a in /](#)

Compressed size: 2.1MB

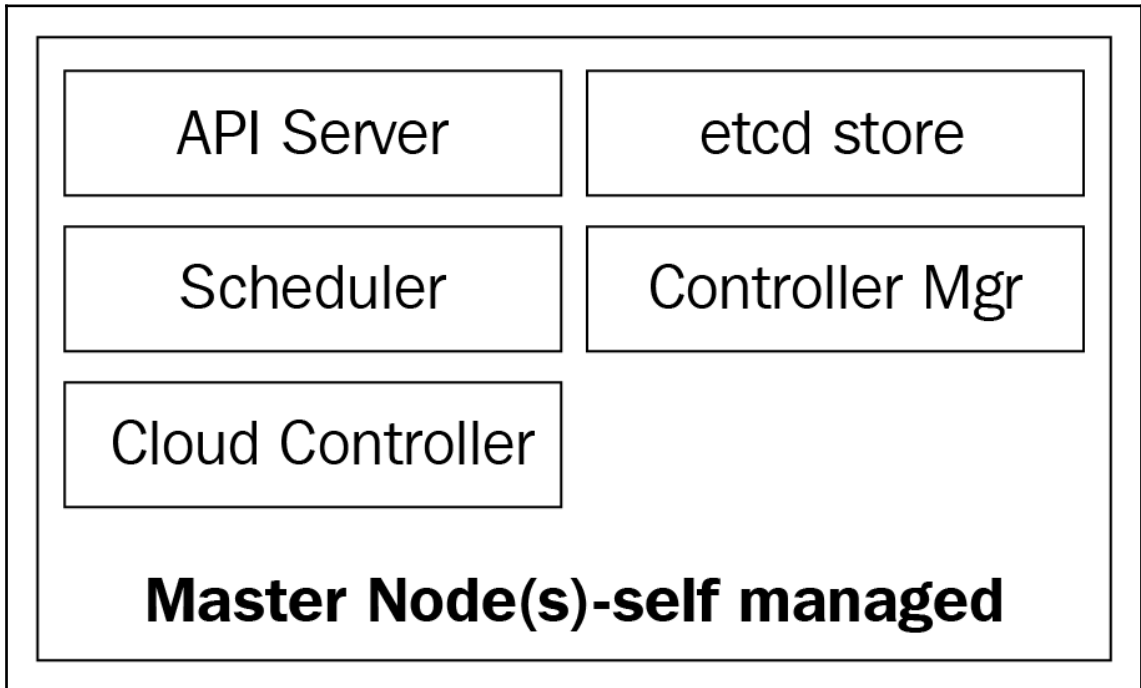


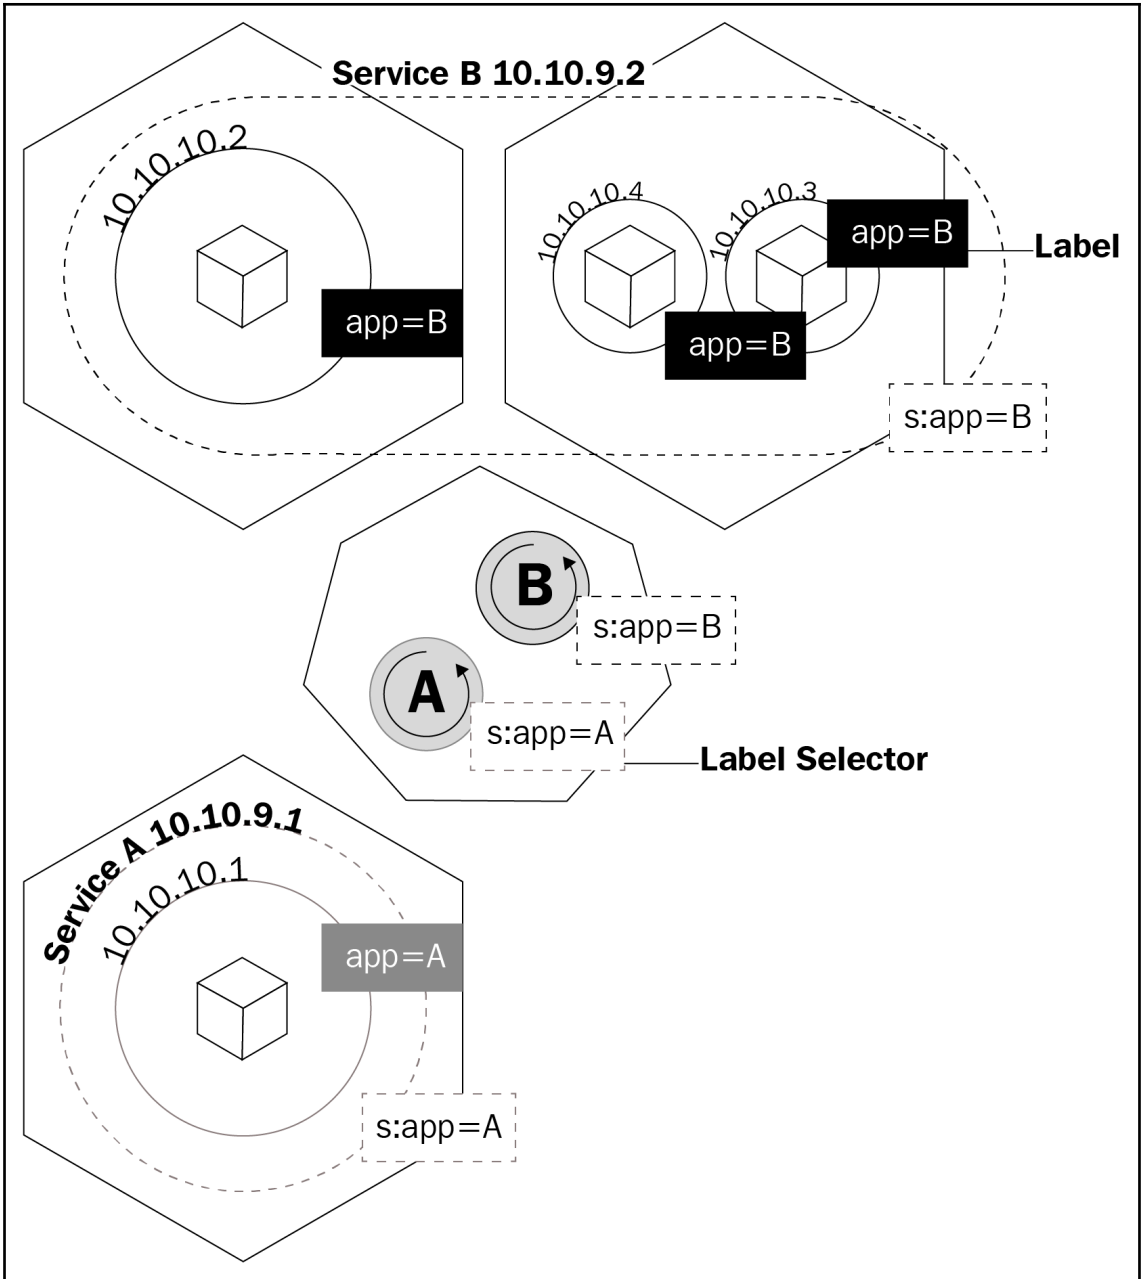
✓ No vulnerable components Base Layer



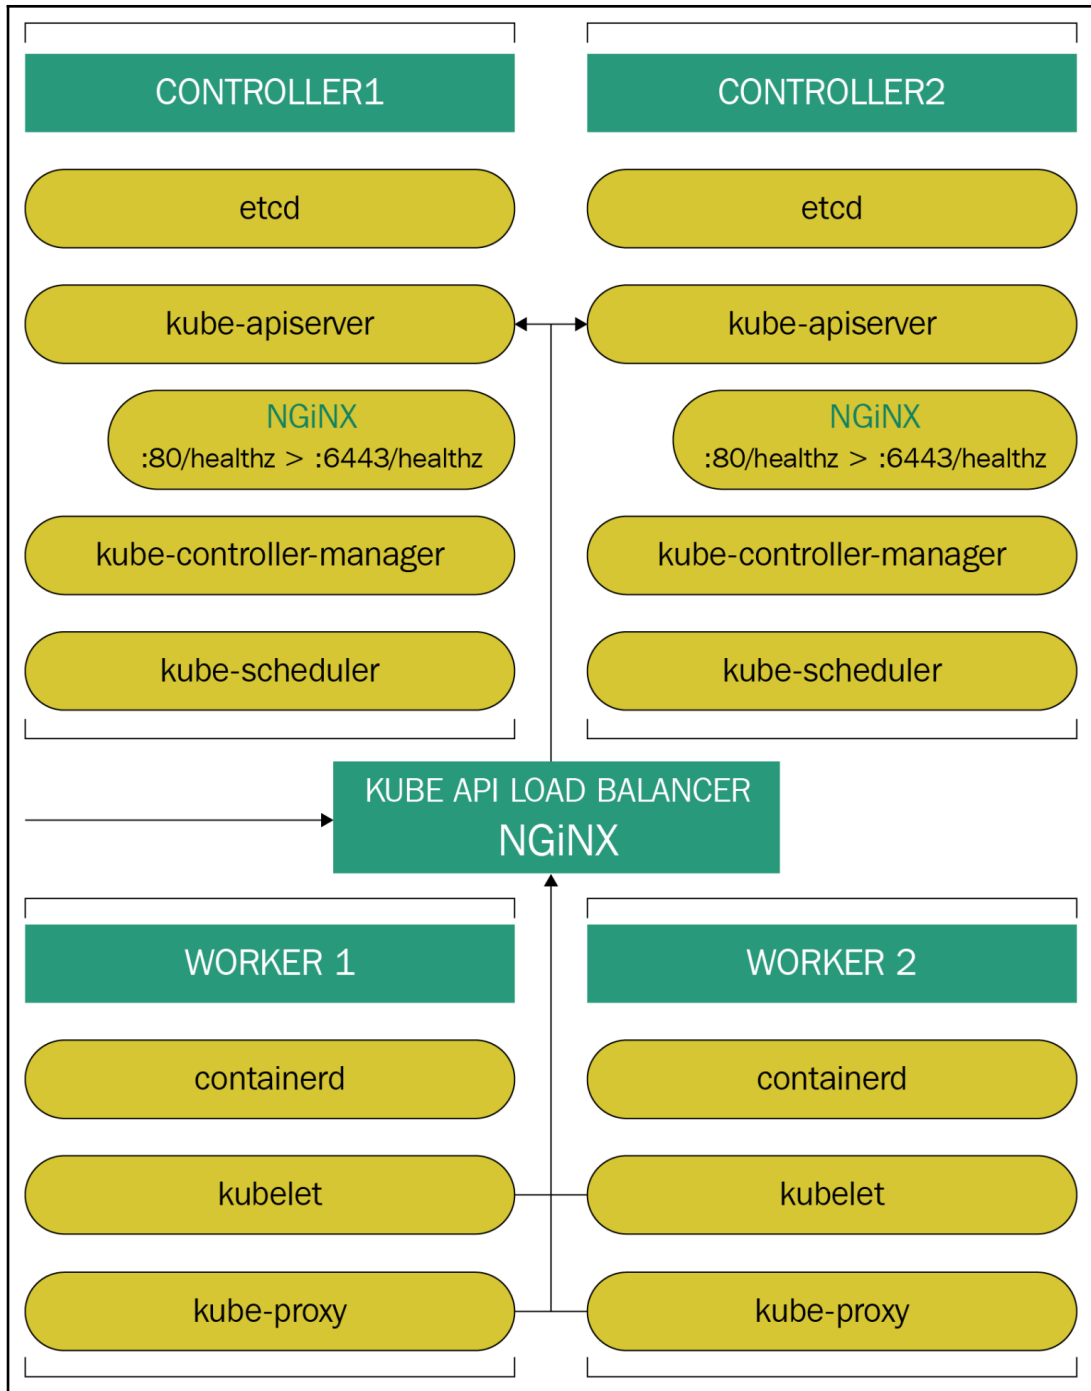


Chapter 7: Understanding the Core Components of a Kubernetes Cluster

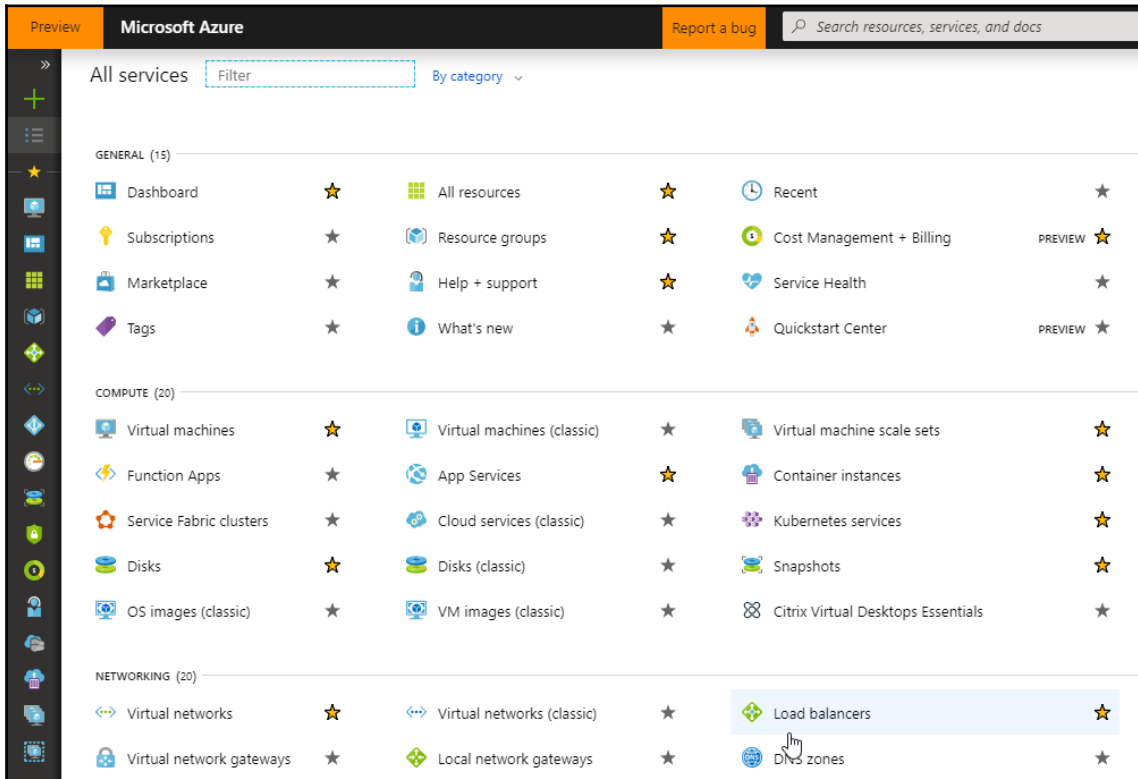


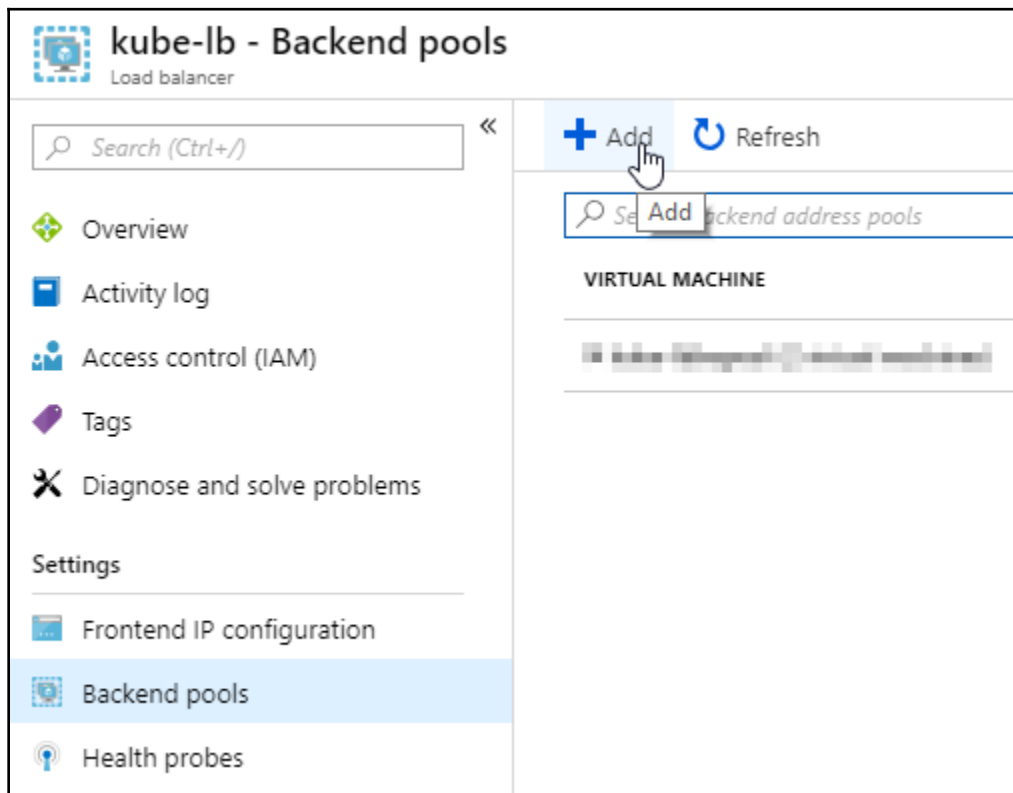
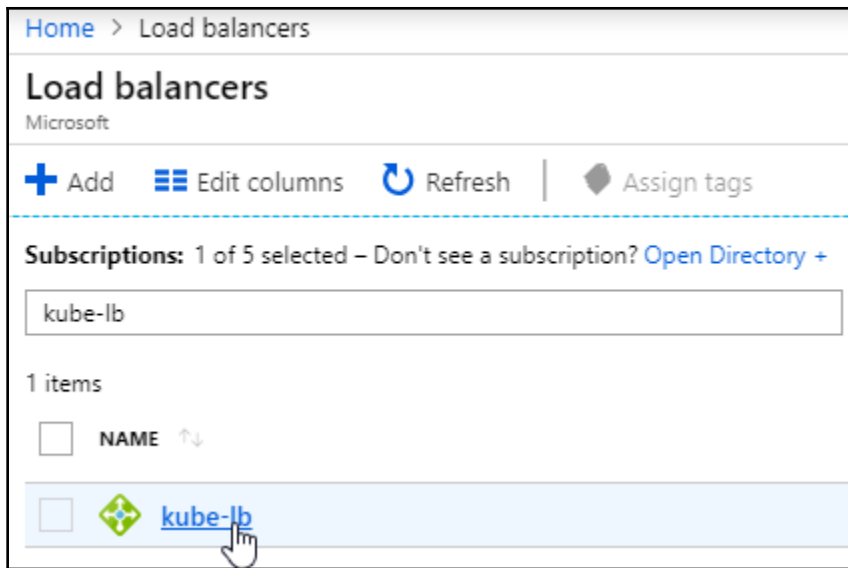


Chapter 8: Architecting a Kubernetes Cluster



Chapter 9: Deploying and Configuring Kubernetes





Add backend pool

kube-lb

* Name
kube-lb-backend ✓

IP version ⓘ
IPv4

Virtual network ⓘ
kube-node-vnet ()

<input checked="" type="checkbox"/>	VIRTUAL MACHINE	IP ADDRESS
<input checked="" type="checkbox"/>	/subscriptions/... ..	/subscriptions/... ..
<input checked="" type="checkbox"/>	/subscriptions/... ..	/subscriptions/... ..
<input checked="" type="checkbox"/>	/subscriptions/... ..	/subscriptions/... ..

Add

Add health probe □ ×

kube-lb

* Name ?

IP version
IPv4

Protocol ?

* Port ?

* Path ?
 ✓

* Interval ?

seconds

* Unhealthy threshold ?

consecutive failures

Add load balancing rule ☐ ✕

kube-lb

*** Name**

kube-api-rule ✓

*** IP Version**

IPv4 IPv6

*** Frontend IP address** ⓘ

104.45.174.96 (LoadBalancerFrontEnd) ▾

Protocol

TCP UDP

*** Port**

6443 ✓

*** Backend port** ⓘ

6443 ✓

Backend pool ⓘ

kube-lbbepool (virtual machines) ▾

Health probe ⓘ

healthz (HTTP:80/healthz) ▾

Session persistence ⓘ

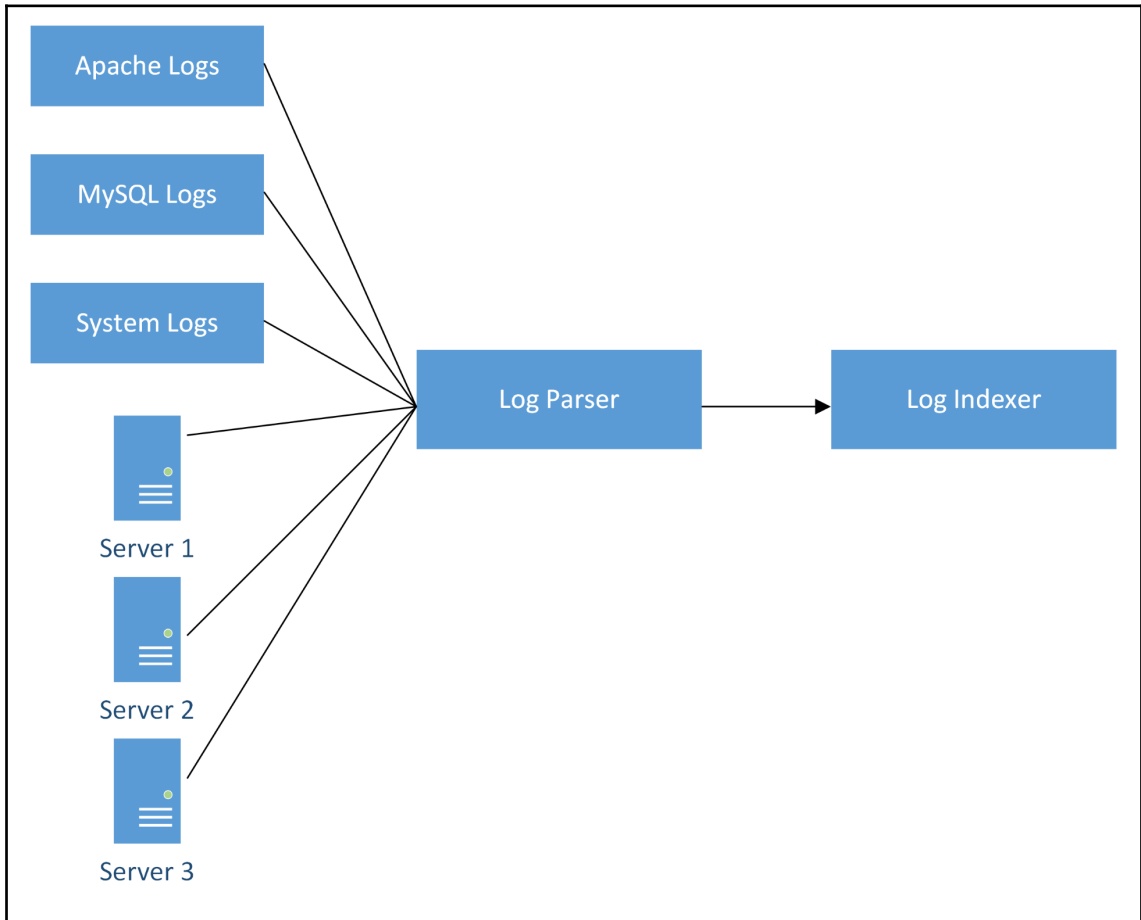
None ▾

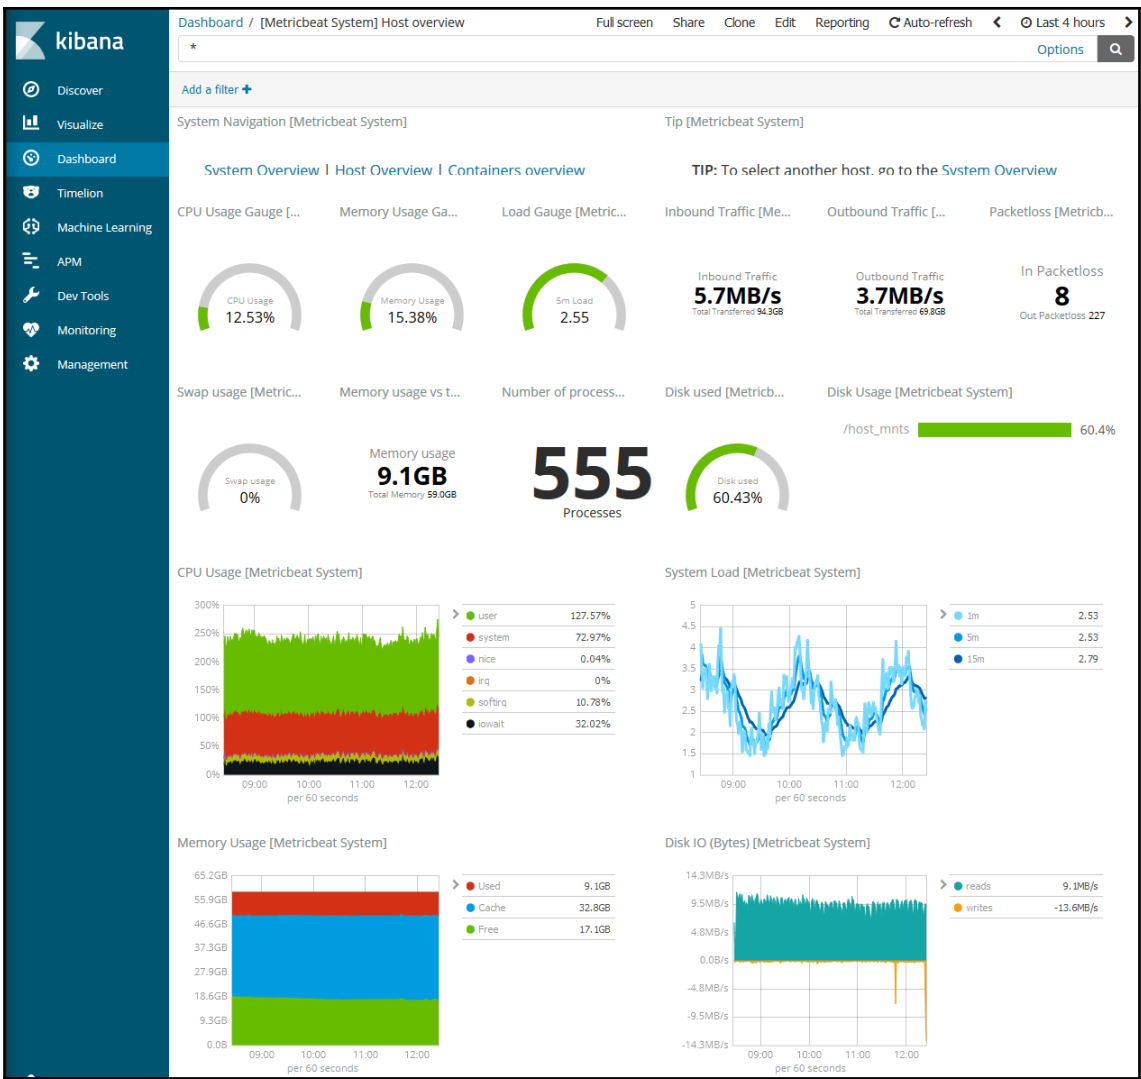
Idle timeout (minutes) ⓘ

4

Floating IP (direct server return) ⓘ

Chapter 10: Monitoring with the ELK Stack



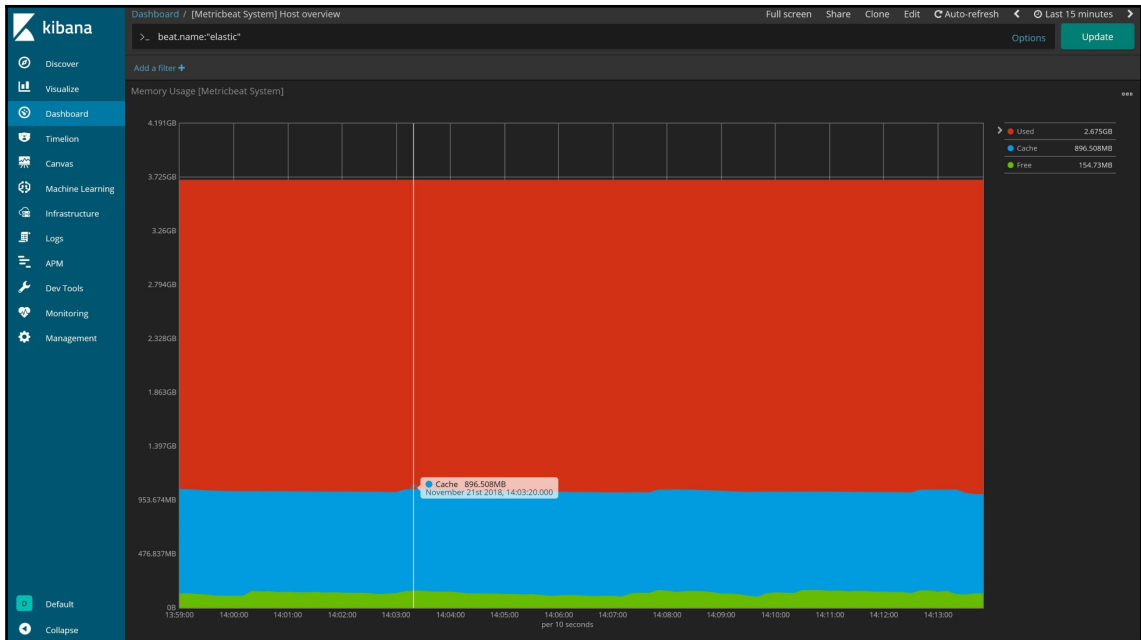


Chapter 11: Designing an ELK Stack



```
[root@elastic ~]# free -m
```

	total	used	free	shared	buff/cache	available
Mem:	3790	2461	103	1	1224	1051
Swap:	2047	20	2027			

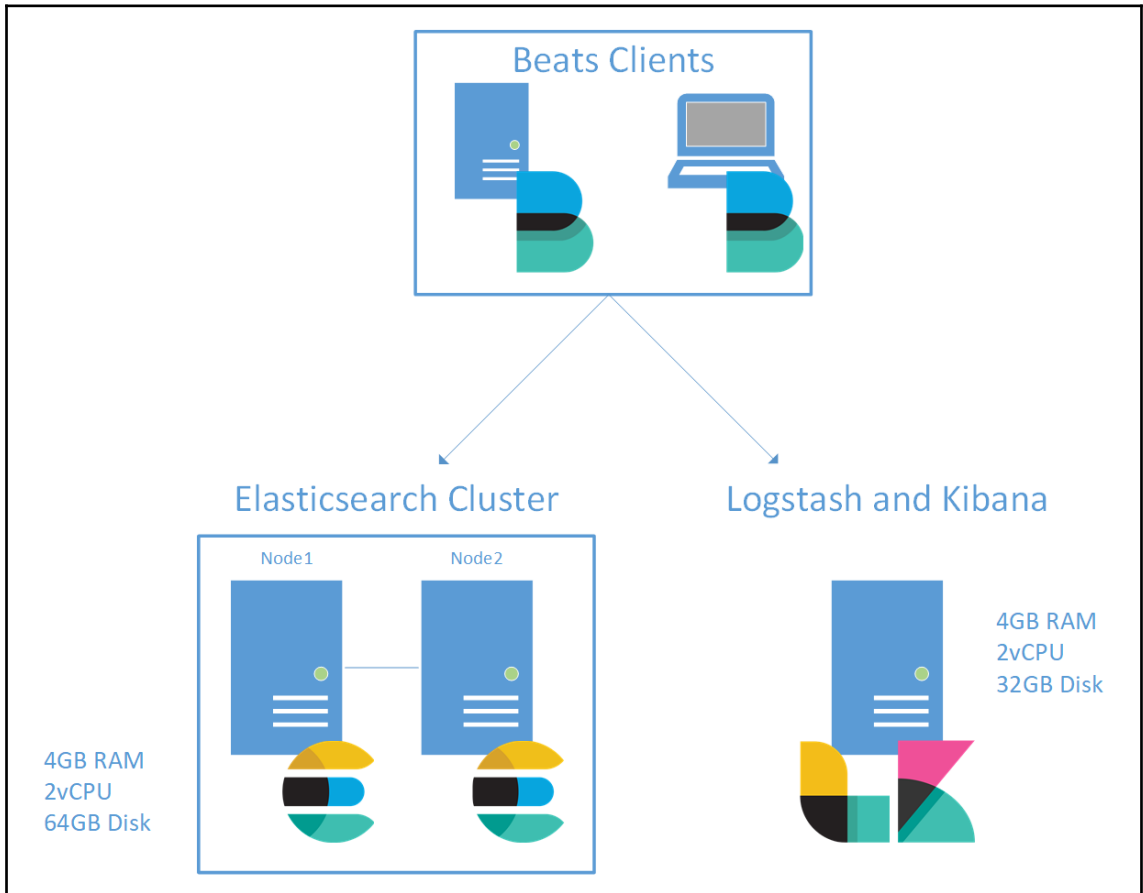


```
[root@elastic ~]# systemctl status elasticsearch -l
● elasticsearch.service - Elasticsearch
   Loaded: loaded (/usr/lib/systemd/system/elasticsearch.service; enabled; vendor preset: disabled)
   Active: failed (Result: exit-code) since Wed 2018-11-21 17:29:05 CST; 2min 42s ago
     Docs: http://www.elastic.co
   Process: 933 ExecStart=/usr/share/elasticsearch/bin/elasticsearch -p ${PID_DIR}/elasticsearch.pid --quiet (code=exited, status=1/FAILURE)
 Main PID: 933 (code=exited, status=1/FAILURE)

Nov 21 17:28:58 elastic systemd[1]: Starting Elasticsearch...
Nov 21 17:29:05 elastic elasticsearch[933]: OpenJDK 64-Bit Server VM warning: INFO: os::commit_memory(0x0000000ca60000, 899284992, 0) failed; error='Cannot allocate memory' (errno=12)
Nov 21 17:29:05 elastic elasticsearch[933]: #
Nov 21 17:29:05 elastic elasticsearch[933]: # There is insufficient memory for the Java Runtime Environment to continue.
Nov 21 17:29:05 elastic elasticsearch[933]: # Native memory allocation (mmap) failed to map 899284992 bytes for committing reserved memory.
Nov 21 17:29:05 elastic elasticsearch[933]: # An error report file with more information is saved as:
Nov 21 17:29:05 elastic elasticsearch[933]: # /var/log/elasticsearch/hs_err_pid933.log
Nov 21 17:29:05 elastic systemd[1]: elasticsearch.service: main process exited, code=exited, status=1/FAILURE
Nov 21 17:29:05 elastic systemd[1]: Unit elasticsearch.service entered failed state.
Nov 21 17:29:05 elastic systemd[1]: elasticsearch.service failed.
```

```
[root@elastic ~]# grep -i "out of memory" /var/log/messages -A4
Nov 21 17:53:20 elastic kernel: Out of memory: Kill process 905 (java) score 576 or sacrifice child
Nov 21 17:53:20 elastic kernel: Killed process 1436 (controller) total-vm:72136kB, anon-rss:648kB, file-rss:0kB, shmem-rss:0kB
Nov 21 17:53:20 elastic kernel: metricbeat invoked oom-killer: gfp_mask=0x201da, order=0, oom_score_adj=0
Nov 21 17:53:20 elastic kernel: metricbeat cpuset=/ mems_allowed=0
Nov 21 17:53:20 elastic kernel: CPU: 1 PID: 917 Comm: metricbeat Kdump: loaded Not tainted 3.10.0-862.14.4.el7.x86_64 #1
...
Nov 21 17:53:20 elastic kernel: Out of memory: Kill process 905 (java) score 576 or sacrifice child
Nov 21 17:53:20 elastic kernel: Killed process 905 (java) total-vm:6944356kB, anon-rss:1376732kB, file-rss:0kB, shmem-rss:0kB
Nov 21 17:53:20 elastic systemd: elasticsearch.service: main process exited, code=killed, status=9/KILL
Nov 21 17:53:20 elastic systemd: Unit elasticsearch.service entered failed state.
Nov 21 17:53:20 elastic systemd: elasticsearch.service failed.
```

Chapter 12: Using Elasticsearch, Logstash, and Kibana to Manage Logs



$$\left(\frac{N}{2}\right) + 1$$

The screenshot displays the Kibana monitoring interface for an 'elastic-cluster'. The dashboard is organized into three main sections: Elasticsearch, Kibana, and Logstash.

Elasticsearch Cluster (Health: yellow)

- Overview:** Version 6.5.4, Uptime 3 minutes.
- Nodes: 3**
 - Disk Available: 97.87% (129.7 GB / 132.5 GB)
 - JVM Heap: 18.00% (544.1 MB / 3.0 GB)
- Indices: 4**
 - Documents: 4,107
 - Disk Usage: 4.2 MB
 - Primary Shards: 4
 - Replica Shards: 2

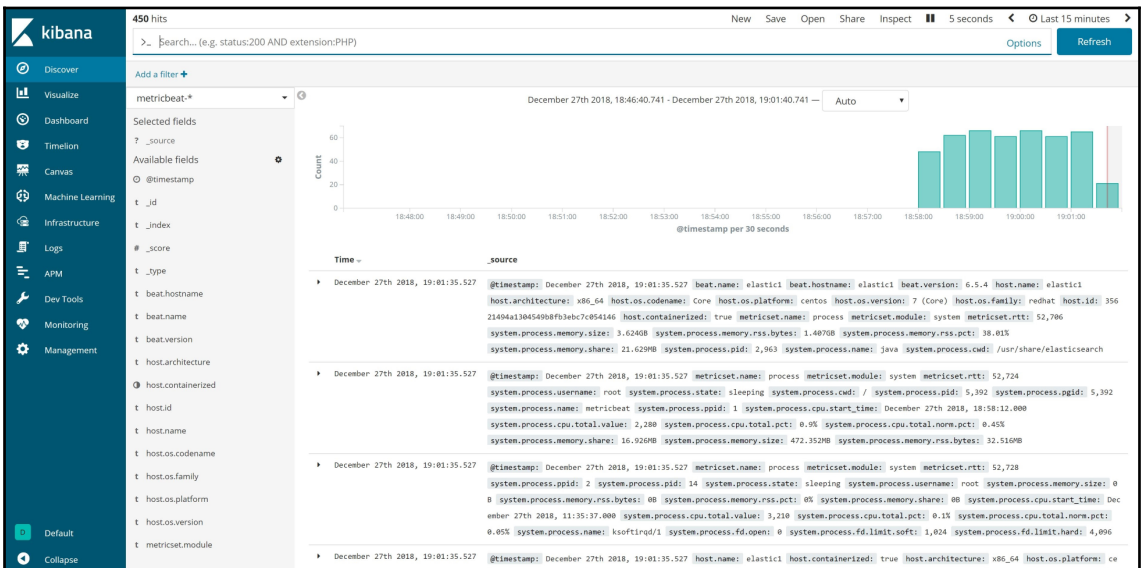
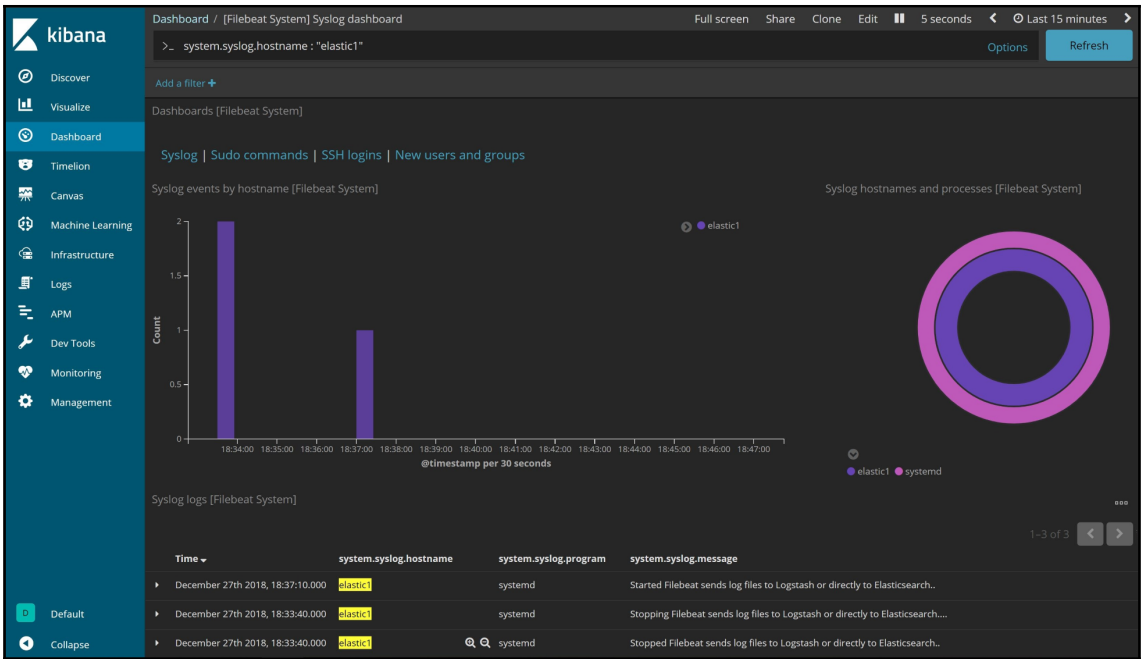
Kibana Cluster (Health: green)

- Overview:** Requests 2, Max. Response Time 149 ms.
- Instances: 1**
 - Connections: 0
 - Memory Usage: 19.62% (281.0 MB / 1.4 GB)

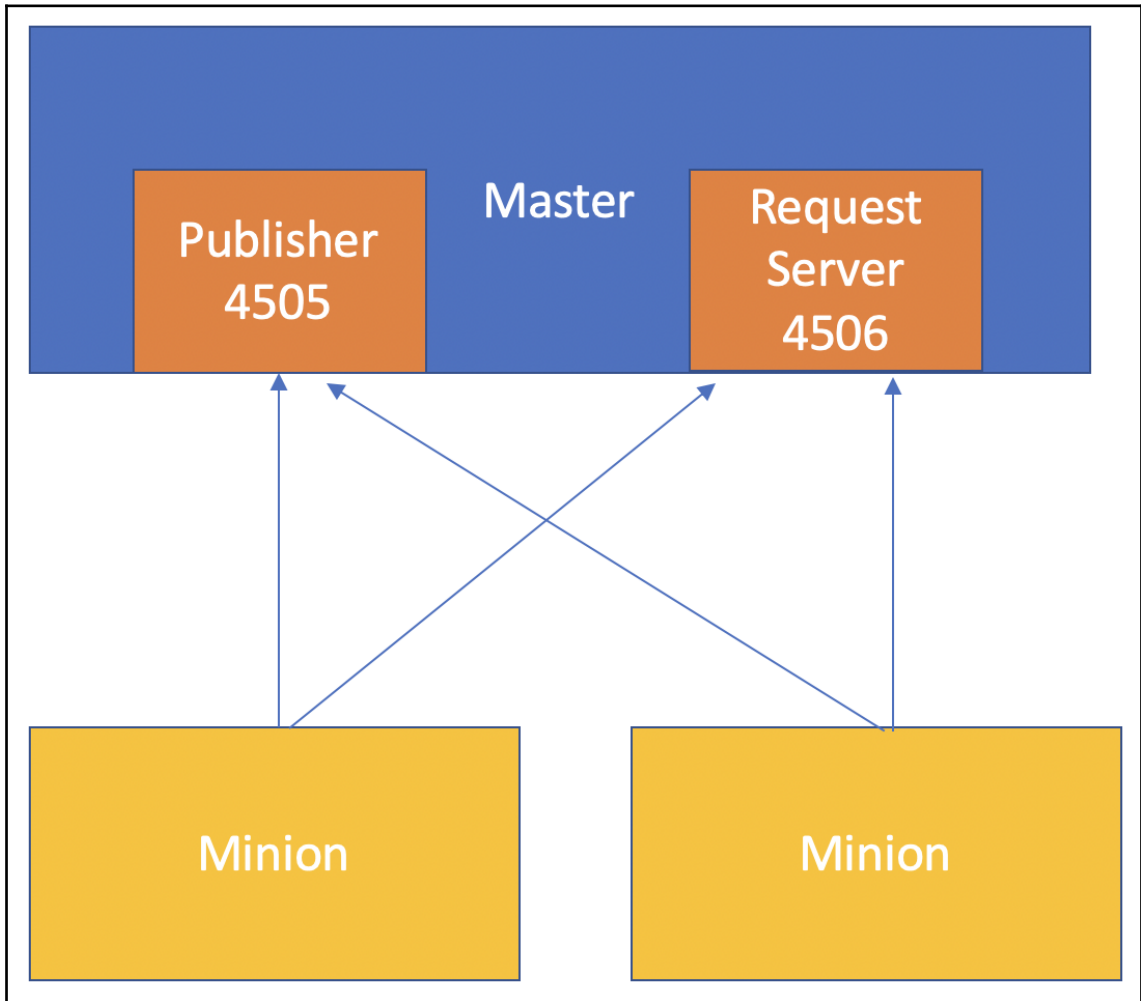
Logstash Cluster

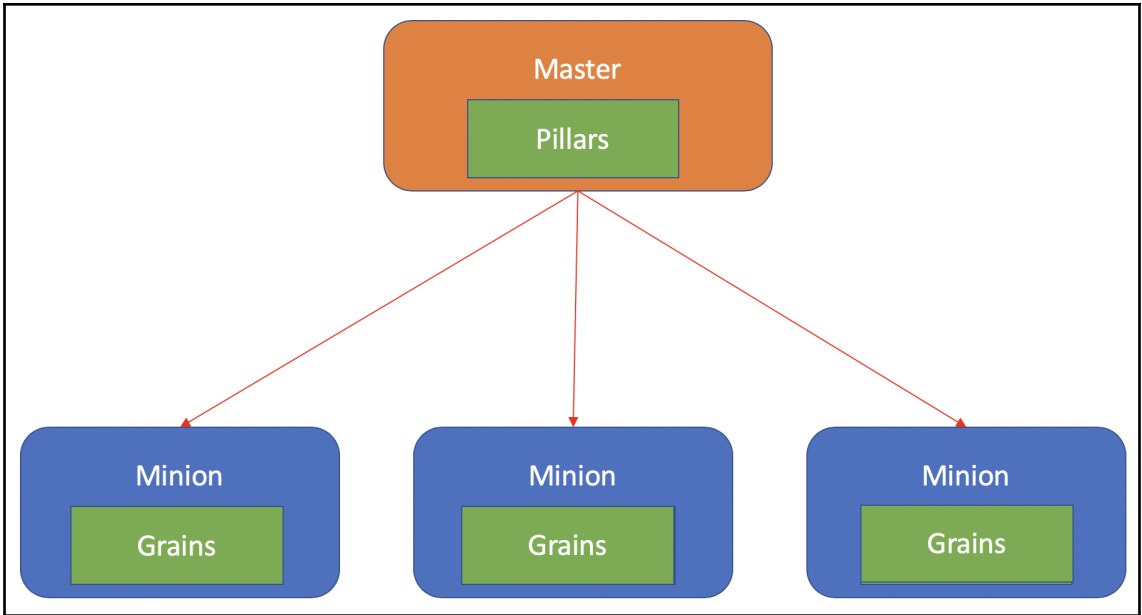
- Overview:** Events Received 0, Events Emitted 0.
- Nodes: 1**
 - Uptime: 3 minutes
 - JVM Heap: 23.40% (235.7 MB / 1,007.4 MB)
- Pipelines: 1**
 - With Memory Queues: 0
 - With Persistent Queues: 1

The left sidebar contains navigation options: Discover, Visualize, Dashboard, Timelion, Canvas, Machine Learning, Infrastructure, Logs, APM, Dev Tools, Monitoring, and Management. The top right shows a refresh rate of 10 seconds and a time range of Last 1 hour.



Chapter 13: Solving Management Problems with Salty Solutions





Chapter 14: Getting Your Hands Salty

```
2481 K8sAzureCNI/default 18:40:15
wget https://releases.hashicorp.com/terraform/0.11.13/terraform_0.11.13_darwin_amd64.zip -O terraform.zip
--2019-04-07 18:42:46-- https://releases.hashicorp.com/terraform/0.11.13/terraform_0.11.13_darwin_amd64.zip
Resolving releases.hashicorp.com (releases.hashicorp.com)... 151.101.5.183
Connecting to releases.hashicorp.com (releases.hashicorp.com)|151.101.5.183|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 22547753 (22M) [application/zip]
Saving to: 'terraform.zip'

terraform.zip          100%[=====>] 21.50M  763KB/s  in 24s
2019-04-07 18:43:11 (902 KB/s) - 'terraform.zip' saved [22547753/22547753]
```

```
1. dsala@NixMachine: ~ (zsh)
2488 K8sAzureCNI/default 18:47:39
unzip terraform.zip -d /usr/local/bin/
Archive: terraform.zip
inflating: /usr/local/bin/terraform
2489 K8sAzureCNI/default 18:47:48
ls -la /usr/local/bin/terraform
-rwxrwxr-x 1 dsala admin 98302152 Mar 12 12:37 /usr/local/bin/terraform
```

```
1. dsala@NixMachine: ~ (zsh)
2491 K8sAzureCNI/default 18:53:42
terraform version
Terraform v0.11.13
2492 K8sAzureCNI/default 18:53:46
```

```
1. dsala@NixMachine: ~/Documents/Code/Salt/terrafiles (zsh)
~/Documents/Code/Salt/terrafiles
└─ terraform init

Initializing provider plugins...
- Checking for available provider plugins on https://releases.hashicorp.com...
- Downloading plugin for provider "azurerm" (1.24.0)...

The following providers do not have any version constraints in configuration,
so the latest version was installed.

To prevent automatic upgrades to new major versions that may contain breaking
changes, it is recommended to add version = "..." constraints to the
corresponding provider blocks in configuration, with the constraint strings
suggested below.

* provider.azurerm: version = "~> 1.24"

Terraform has been successfully initialized!

You may now begin working with Terraform. Try running "terraform plan" to see
any changes that are required for your infrastructure. All Terraform commands
should now work.

If you ever set or change modules or backend configuration for Terraform,
rerun this command to reinitialize your working directory. If you forget, other
commands will detect it and remind you to do so if necessary.
~/Documents/Code/Salt/terrafiles
```

```
1. dsala@NixMachine: ~/Documents/Code/Salt/terrafiles (zsh)
~/Documents/Code/Salt/terrafiles
└─ terraform plan

Warning: azurerm_public_ip.saltnginxpip: "public_ip_address_allocation": [DEPRECATED] this property has been
removed in the API

Refreshing Terraform state in-memory prior to plan...
The refreshed state will be used to calculate this plan, but will not be
persisted to local or remote state storage.

-----

An execution plan has been generated and is shown below.
Resource actions are indicated with the following symbols:
+ create

Terraform will perform the following actions:

+ azurerm_network_interface.saltlb
  id:                                     <computed>
  applied_dns_servers.#:                 <computed>
  dns_servers.#:                         <computed>
  enable_accelerated_networking:         "false"
  enable_ip_forwarding:                  "false"
  internal_dns_name_label:               <computed>
```

```
1. terraform apply (terraform-provid)

storage_os_disk.0.caching:                "ReadWrite"
storage_os_disk.0.create_option:          "FromImage"
storage_os_disk.0.disk_size_gb:           <computed>
storage_os_disk.0.managed_disk_id:        <computed>
storage_os_disk.0.managed_disk_type:      "Standard_LRS"
storage_os_disk.0.name:                   "webosdisk1"
storage_os_disk.0.write_accelerator_enabled: "false"
tags.%:                                    <computed>
vm_size:                                   "Standard_B1s"

+ azurerm_virtual_network.salt
  id:                                       <computed>
  address_space.#:                          "1"
  address_space.0:                          "10.0.0.0/16"
  location:                                  "eastus"
  name:                                       "saltnet"
  resource_group_name:                       "Salt"
  subnet.#:                                  <computed>
  tags.%:                                    <computed>

Plan: 13 to add, 0 to change, 0 to destroy.

Do you want to perform these actions?
  Terraform will perform the actions described above.
  Only 'yes' will be accepted to approve.

Enter a value: █
```



```
1. dsala@salt:~ (ssh)
[dsala@salt ~]$ systemctl status salt-master
● salt-master.service - The Salt Master Server
  Loaded: loaded (/usr/lib/systemd/system/salt-master.service; enabled; vendor preset: disabled)
  Active: active (running) since Mon 2019-04-08 17:19:02 UTC; 8s ago
  Docs: man:salt-master(1)
        file:///usr/share/doc/salt/html/contents.html
        https://docs.saltstack.com/en/latest/contents.html
  Main PID: 66418 (salt-master)
  CGroup: /system.slice/salt-master.service
          └─66418 /usr/bin/python /usr/bin/salt-master
            └─66432 /usr/bin/python /usr/bin/salt-master
              └─66437 /usr/bin/python /usr/bin/salt-master
                └─66438 /usr/bin/python /usr/bin/salt-master
                  └─66441 /usr/bin/python /usr/bin/salt-master
                    └─66442 /usr/bin/python /usr/bin/salt-master
                      └─66443 /usr/bin/python /usr/bin/salt-master
                        └─66444 /usr/bin/python /usr/bin/salt-master
                          └─66445 /usr/bin/python /usr/bin/salt-master
                            └─66446 /usr/bin/python /usr/bin/salt-master
                              └─66453 /usr/bin/python /usr/bin/salt-master
                                └─66454 /usr/bin/python /usr/bin/salt-master
                                  └─66455 /usr/bin/python /usr/bin/salt-master
                                    └─66544 /usr/bin/python /usr/bin/salt-master
                                      └─66545 /usr/bin/python /usr/bin/salt-master
                                        └─66546 /usr/bin/python /usr/bin/salt-master
                                          └─66547 /usr/bin/python /usr/bin/salt-master
                                            └─66549 /usr/bin/python /usr/bin/salt-master
[dsala@salt ~]$
```

```
1. dsala@lb-vm:~ (ssh)
dsala@salt:~ (ssh) #1  dsala@lb-vm:~ (ssh) #2
[dsala@lb-vm ~]$ sudo systemctl status salt-minion
● salt-minion.service - The Salt Minion
   Loaded: loaded (/usr/lib/systemd/system/salt-minion.service; enabled; vendor preset: disabled)
   Active: active (running) since Mon 2019-04-08 19:08:43 UTC; 4min 49s ago
     Docs: man:salt-minion(1)
           file:///usr/share/doc/salt/html/contents.html
           https://docs.saltstack.com/en/latest/contents.html
   Main PID: 71779 (salt-minion)
   CGroup: /system.slice/salt-minion.service
           └─71779 /usr/bin/python /usr/bin/salt-minion
             └─71782 /usr/bin/python /usr/bin/salt-minion
               └─71790 /usr/bin/python /usr/bin/salt-minion

Apr 08 19:11:55 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Apr 08 19:12:05 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Apr 08 19:12:15 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Apr 08 19:12:25 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Apr 08 19:12:35 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Apr 08 19:12:45 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Apr 08 19:12:55 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Apr 08 19:13:05 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Apr 08 19:13:15 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Apr 08 19:13:25 lb-vm salt-minion[71779]: [ERROR ] The Salt Master has cached the public key for thi...cate
Hint: Some lines were ellipsized, use -l to show in full.
[dsala@lb-vm ~]$
```

```
1. dsala@web-00: ~ (ssh)
dsala@web-00:~ (ssh) #1  ~ (zsh) #2
dsala@web-00:~$ sudo wget -O - https://repo.saltstack.com/apt/ubuntu/16.04/amd64/latest/SALTSTACK-GPG-KEY.pub
| sudo apt-key add -
--2019-04-08 19:21:25-- https://repo.saltstack.com/apt/ubuntu/16.04/amd64/latest/SALTSTACK-GPG-KEY.pub
Resolving repo.saltstack.com (repo.saltstack.com)... 138.197.226.47, 2604:a880:400:d0::2:e001
Connecting to repo.saltstack.com (repo.saltstack.com)|138.197.226.47|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 1727 (1.7K)
Saving to: 'STDOUT'

-
100%[=====>] 1.69K --.-KB/s in 0s

2019-04-08 19:21:25 (27.7 MB/s) - written to stdout [1727/1727]

OK
dsala@web-00:~$
```

```
1. dsala@web-00: ~ (ssh)
dsala@web-00:~$ sudo apt update
Hit:1 http://azure.archive.ubuntu.com/ubuntu xenial InRelease
Hit:2 http://azure.archive.ubuntu.com/ubuntu xenial-updates InRelease
Get:3 http://azure.archive.ubuntu.com/ubuntu xenial-backports InRelease [107 kB]
Hit:4 http://repo.saltstack.com/apt/ubuntu/16.04/amd64/latest xenial InRelease ←
Get:5 http://security.ubuntu.com/ubuntu xenial-security InRelease [109 kB]
Fetched 216 kB in 0s (463 kB/s)
Reading package lists... Done
Building dependency tree
Reading state information... Done
11 packages can be upgraded. Run 'apt list --upgradable' to see them.
dsala@web-00:~$
```

```
1. dsala@salt:~ (ssh)
..lt/terraformfiles (zsh) 1 x ~ (zsh) 2 x dsala@salt:~ (ssh) 3
#####
# This configuration file is used to manage the behavior of the Salt Master.
# Values that are commented out but have an empty line after the comment are
# defaults that do not need to be set in the config. If there is no blank line
# after the comment then the value is presented as an example and is not the
# default.

# Per default, the master will automatically include all config files
# from master.d/*.conf (master.d is a directory in the same directory
# as the main master config file).
#default_include: master.d/*.conf

# The address of the interface to bind to:
#interface: 0.0.0.0
interface: 10.0.0.10

# Whether the master should listen for IPv6 connections. If this is set to True,
# the interface option must be adjusted, too. (For example: "interface: ':::')
#ipv6: False

# The tcp port used by the publisher:
#publish_port: 4505

# The user under which the salt master will run. Salt will update all
# permissions to allow the specified user to run the master. The exception is
# the job cache, which must be deleted if this user is changed. If the
27,1 0%
```

```
1. dsala@salt:~ (ssh)
x ..lt/terraformfiles (zsh) #1 x ~ (zsh) #2 x dsala@salt:~ (ssh) #3
[dsala@salt ~]$ ss -tunap | grep 450
tcp    LISTEN    0      128     10.0.0.10:4505      *:*
tcp    LISTEN    0      128     10.0.0.10:4506      *:*
```

```
1. dsala@web-00: ~ (ssh)
x ..lt/terraformfiles (zsh) #1 x ~ (zsh) #2 x dsala@web-00: ~ (... #3
# not be set in the config. If there is no blank line after the comment, the
# value is presented as an example and is not the default.

# Per default the minion will automatically include all config files
# from minion.d/*.conf (minion.d is a directory in the same directory
# as the main minion config file).
#default_include: minion.d/*.conf

# Set the location of the salt master server. If the master server cannot be
# resolved, then the minion will fail to start.
#master: salt
master: 10.0.0.10

# Set http proxy information for the minion when doing requests
```

```
1. dsala@salt:~ (ssh)
x dsala@salt:~ (ssh) #1 x ...a@web-00: ~ (ssh) #2
[dsala@salt ~]$ sudo salt-key -f master.pub
Local Keys:
master.pub: 29:2d:15:09:ab:e0:7d:dc:09:d7:9b:d6:a8:04:50:48:b9:11:ee:3b:a0:3c:b0:68:0b:37:f4:9e:92:a7:8a:55
[dsala@salt ~]$
```

```
1. dsala@web-00: ~ (ssh)
dsala@salt:~ (ssh) 1
...a@web-00: ~ (ssh) 2
# The state_output_diff setting changes whether or not the output from
# successful states is returned. Useful when even the terse output of these
# states is cluttering the logs. Set it to True to ignore them.
#state_output_diff: False

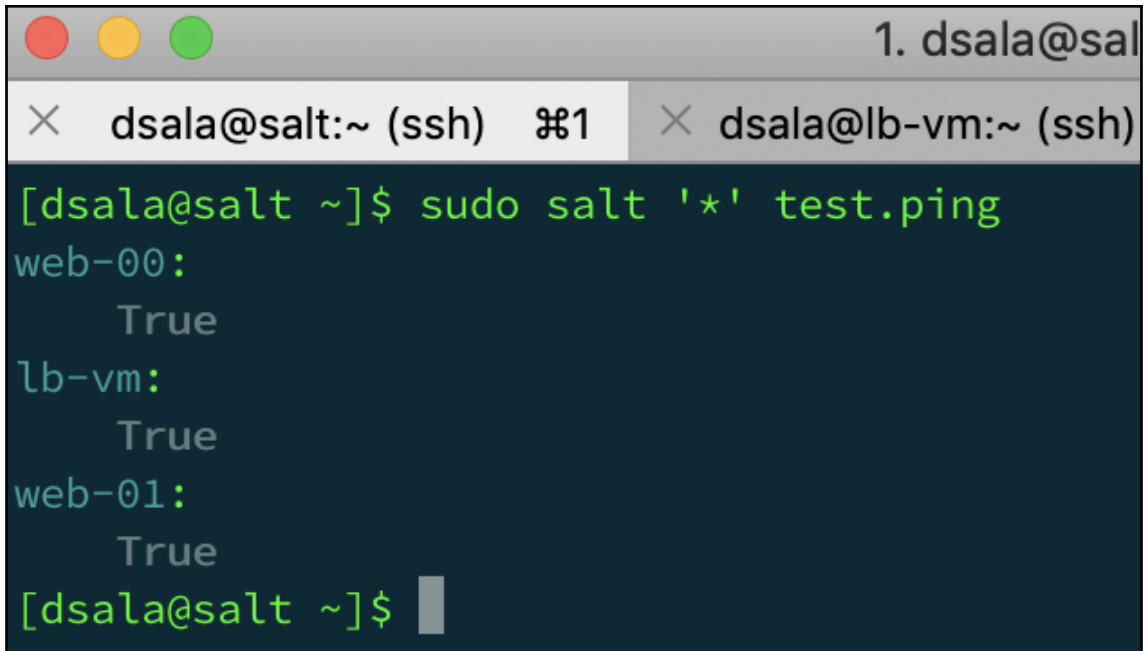
# The state_output_profile setting changes whether profile information
# will be shown for each state run.
#state_output_profile: True

# Fingerprint of the master public key to validate the identity of your Salt master
# before the initial key exchange. The master fingerprint can be found by running
# "salt-key -f master.pub" on the Salt master.
master_finger: '29:2d:15:09:ab:e0:7d:dc:09:d7:9b:d6:a8:04:50:48:b9:11:ee:3b:a0:3c:b0:68:0b:37:f4:9e:92:a7:8a:55'
```

```
1. dsala@salt:~ (ssh)
dsala@salt:~ (ssh) 1
dsala@lb-vm:~ (ssh) 2
[dsala@salt ~]$ sudo salt-key -F
Local Keys:
master.pem: 17:6f:4a:8d:47:60:b2:c9:46:42:1d:10:5c:91:96:a5:69:18:6e:de:d3:d6:8f:d4:d8:9c:58:73:7f:1c:c4:33
master.pub: 29:2d:15:09:ab:e0:7d:dc:09:d7:9b:d6:a8:04:50:48:b9:11:ee:3b:a0:3c:b0:68:0b:37:f4:9e:92:a7:8a:55
Unaccepted Keys:
lb-vm: c1:dc:8c:e4:58:4d:cb:3b:99:b9:b5:51:0d:5c:3f:50:ff:bc:20:65:3f:76:22:1c:96:10:b6:59:94:a0:e9:ac
web-00: f3:96:5e:63:f6:63:5a:c3:fa:1a:b1:ab:8b:c4:e0:0c:73:bc:1f:14:0c:a5:3a:65:01:dc:4d:19:c7:24:52:59
web-01: 38:31:c7:e8:a3:a2:3c:81:1d:9d:ad:dc:78:5a:29:ec:fc:23:98:8c:08:21:f7:10:31:5a:fb:8e:d8:89:27:34

1. dsala@lb-vm:~ (ssh)
dsala@salt:~ (ssh) 1
dsala@lb-vm:~ (ssh) 2
[dsala@lb-vm ~]$ sudo salt-call --local key.finger
local:
c1:dc:8c:e4:58:4d:cb:3b:99:b9:b5:51:0d:5c:3f:50:ff:bc:20:65:3f:76:22:1c:96:10:b6:59:94:a0:e9:ac
[dsala@lb-vm ~]$
```

```
1. dsala@salt
dsala@salt:~ (ssh)  #1  dsala@lb-vm:~ (ssh)
[dsala@salt ~]$ sudo salt-key -A
The following keys are going to be accepted:
Unaccepted Keys:
lb-vm
web-00
web-01
Proceed? [n/Y] Y
Key for minion lb-vm accepted.
Key for minion web-00 accepted.
Key for minion web-01 accepted.
[dsala@salt ~]$ sudo salt-key -L
Accepted Keys:
lb-vm
web-00
web-01
Denied Keys:
Unaccepted Keys:
Rejected Keys:
[dsala@salt ~]$
```



```
1. dsala@sal
dsala@salt:~ (ssh) #1 dsala@lb-vm:~ (ssh)
[dsala@salt ~]$ sudo salt '*' test.ping
web-00:
    True
lb-vm:
    True
web-01:
    True
[dsala@salt ~]$
```

```
Summary for web-01
```

```
-----
```

```
Succeeded: 2 (changed=1)
```

```
Failed:    0
```

```
-----
```

```
Total states run:    2
```

```
Total run time:    70.596 s
```

```
[dsala@salt salt]$
```

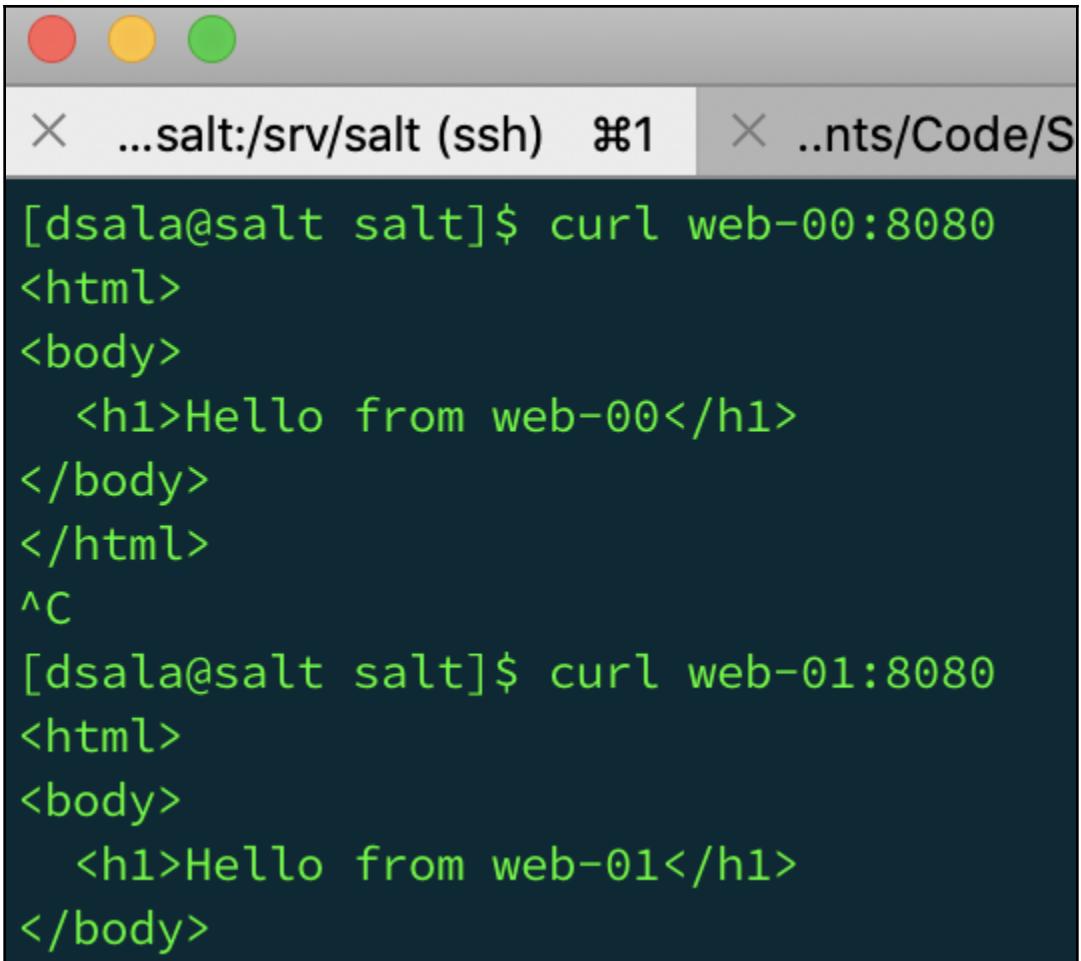


```
-----
      ID: webuser
Function: user.present
      Result: True
      Comment: New user webuser created
      Started: 05:44:54.292257
Duration: 506.713 ms
Changes:
-----
      fullname:
      gid:
          4000
      groups:
          - webuser
      home:
          /home/webuser
      homephone:
      name:
          webuser
      other:
      passwd:
          x
      roomnumber:
      shell:
      uid:
          4000
      workphone:

Summary for web-01
-----
Succeeded: 3 (changed=1)
Failed:    0
-----
Total states run:    3
Total run time: 586.187 ms
[dsala@salt salt]$ █ 0 Azure: desalama
```

```
-----
ID: node-app
Function: git.latest
  Name: https://github.com/dsalamancaMS/SaltChap.git
  Result: True
  Comment: https://github.com/dsalamancaMS/SaltChap.git cloned to /home/webuser/app
  Started: 06:12:50.037084
  Duration: 2557.135 ms
  Changes:
    -----
    new:
      https://github.com/dsalamancaMS/SaltChap.git => /home/webuser/app
    revision:
      -----
      new:
        7726a5c4b30e20fc26448884079f760d833a0611
      old:
        None

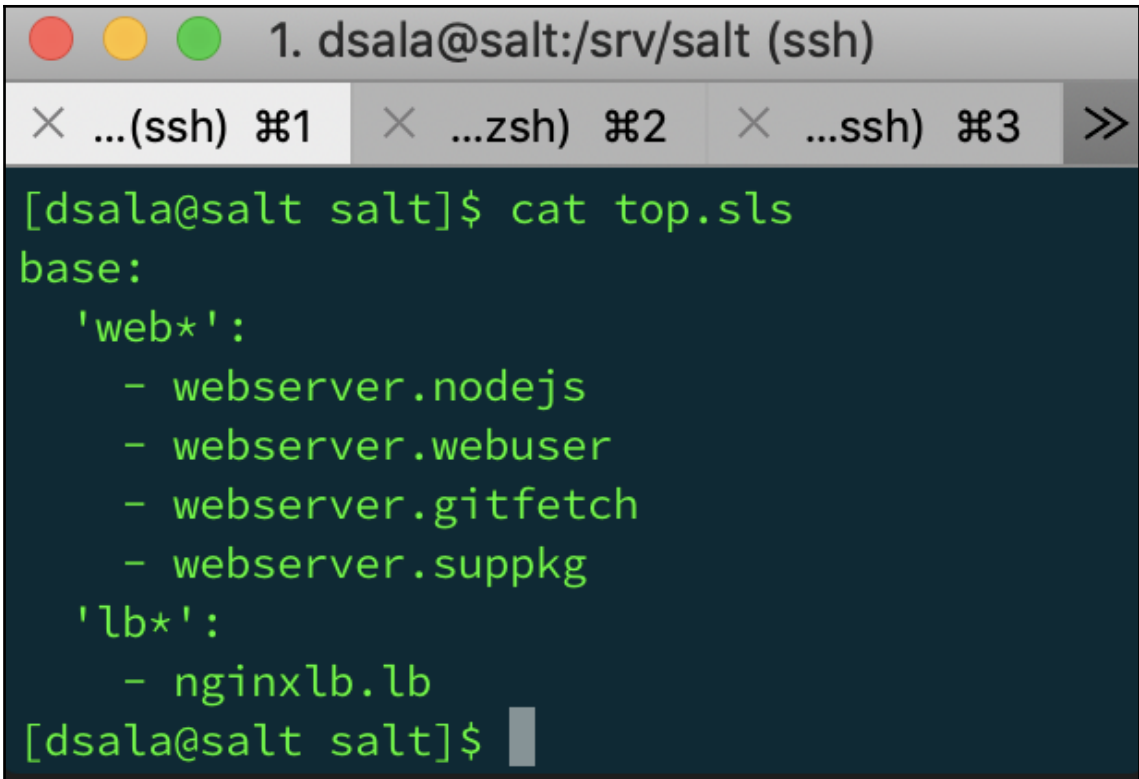
Summary for web-01
-----
Succeeded: 4 (changed=1)
Failed:    0
-----
Total states run:    4
Total run time:    2.642 s
[dsala@salt salt]$ █
```



```
[dsala@salt salt]$ curl web-00:8080
<html>
<body>
  <h1>Hello from web-00</h1>
</body>
</html>
^C
[dsala@salt salt]$ curl web-01:8080
<html>
<body>
  <h1>Hello from web-01</h1>
</body>
```

```
1. dsala@salt:/srv/salt (ssh)
× ...alt (ssh) ⌘1 × ...lt (zsh) ⌘2 × ...p (ssh) ⌘3 × ...:~ (ssh) ⌘4
[dsala@salt salt]$ sudo salt 'web-00' grains.set roles frontend
web-00:
-----
changes:
-----
  roles:
    frontend
comment:
result:
  True
[dsala@salt salt]$
```

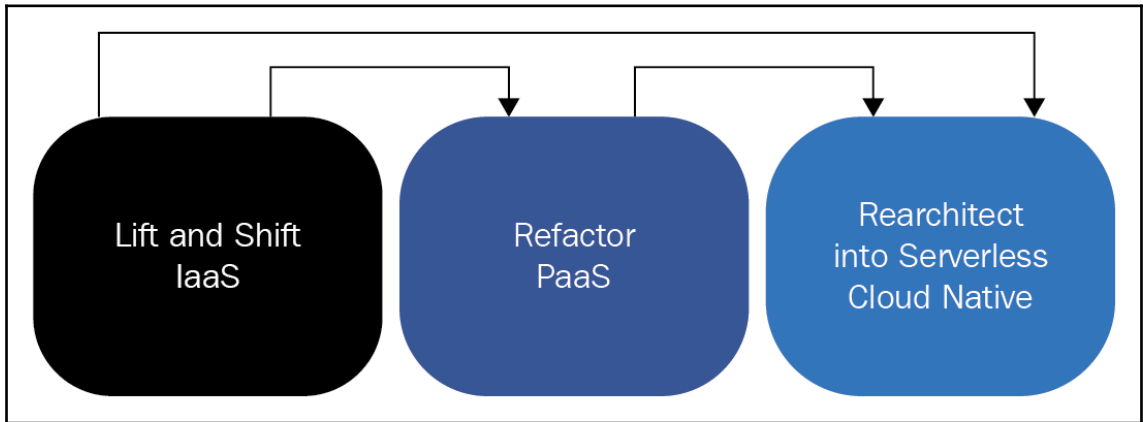
```
[dsala@salt salt]$ sudo salt 'web-01' grains.items | grep roles -A1
roles:
  frontend
```

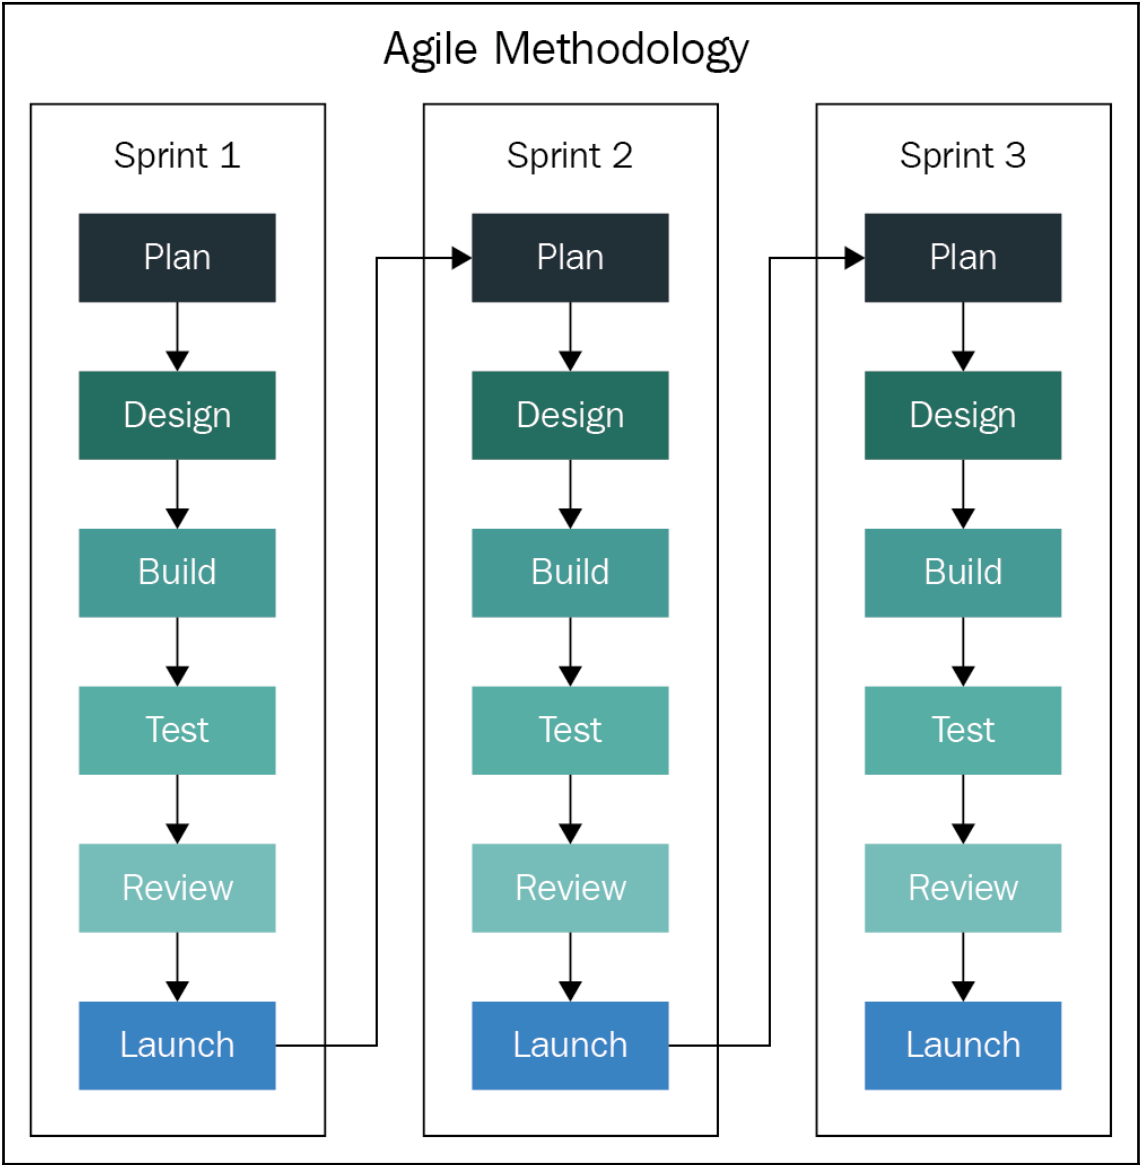


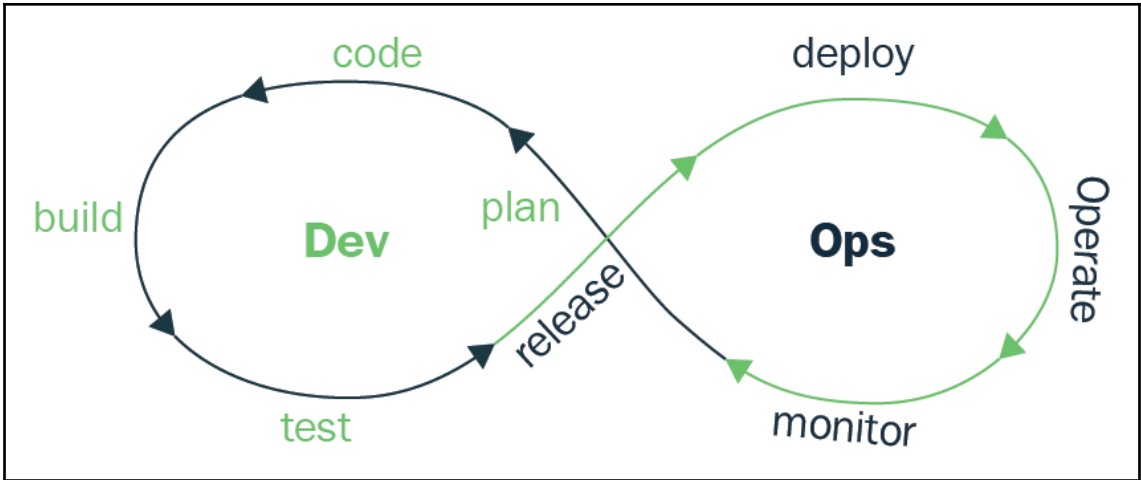
```
1. dsala@salt:/srv/salt (ssh)
× ...(ssh) ⌘1 × ...zsh) ⌘2 × ...ssh) ⌘3 >>
[dsala@salt salt]$ cat top.sls
base:
  'web*':
    - webserver.nodejs
    - webserver.webuser
    - webserver.gitfetch
    - webserver.suppkg
  'lb*':
    - nginxlb.lb
[dsala@salt salt]$
```

```
1. dsala@NixMa
dsala@salt:/srv/...  ⌘1  ..nts/Code/Salt (zsh) ⌘2
~/Documents/Code/Salt  master
curl http://168.61.33.201:8080/
<html>
<body>
  <h1>Hello Reader from web-01</h1>
</body>
</html>
curl: (18) transfer closed with 5 bytes remaining to read
~/Documents/Code/Salt  master
curl http://168.61.33.201:8080/
<html>
<body>
  <h1>Hello Reader from web-00</h1>
</body>
</html>
curl: (18) transfer closed with 5 bytes remaining to read
```

Chapter 15: Design Best Practices







Graphics Bundle Ends Here

Index