
Container Networking

Michael Hausenblas

Beijing • Boston • Farnham • Sebastopol • Tokyo

O'REILLY®

Table of Contents

Preface.....	v
1. Motivation.....	1
Introducing Pets vs. Cattle	1
Go Cattle!	2
Container Networking Stack	3
Do I Need to Go “All In”?	4
2. Introduction to Container Networking.....	7
Single-host Container Networking 101	7
Bridge Mode Networking	9
Host Mode Networking	10
Container Mode Networking	12
No Networking	12
Administrative Considerations	13
Wrapping It Up	14
3. Multi-host networking.....	15
Multi-host Container Networking 101	15
Flannel	15
Weave Net	16
Project Calico	16
Open vSwitch	17
OpenVPN	17
Docker Networking	17
Administrative Considerations	18
Wrapping It Up	19

4. Orchestration.....	21
What Does a Scheduler Actually Do?	22
Docker	24
Apache Mesos	26
Hashicorp Nomad	28
Community Matters	30
Wrapping It Up	31
5. Service Discovery.....	33
The Challenge	33
Technologies	35
Load Balancing	41
Wrapping It Up	42
6. The Container Network Interface.....	45
History	46
Specification and Usage	46
Container Runtimes and Plug-ins	48
Wrapping It Up	50
7. Kubernetes Networking.....	51
A Gentle Kubernetes Introduction	51
Kubernetes Networking Overview	53
Intra-pod networking	55
Inter-pod networking	56
Service Discovery In Kubernetes	59
Ingress And Egress	62
Advanced Kubernetes Networking Topics	65
Wrapping It Up	67
A. References.....	69
Index.....	73

Index

A

Ansible, 3
Apache Cassandra, 26
Apache Kafka, 9
Apache Mesos, 26
Apache Spark, 9, 26
Apache ZooKeeper, 35

B

base provisioning, 22
Border Gateway Protocol, 16
bridge mode, 9
Buoyant, 67

C

CAP theorem, 35
cattle approach
 challenges, 2
Chronos, 26
cluster, 8
CNCF, 46
CNI
 architecture, 47
 overview, 45
 plug-ins, 48
commodity hardware, 2
community, 30
Conduit, 67, 67
Consul, 38
container image, 23
container mode, 12
Container networking stack, 3

container registry, 23
CoreDNS, 61
CRI-O, 7

D

data locality, 9
DC/OS, 28
distributed system, 4
 scheduler, 22
DNS, 60
Docker, 7, 9
Docker security, 14

E

elasticity, 2
Envoy, 41
etcd, 37
Eureka, 40

F

Facebook, 8

G

Google
 Borg, 22, 51
 Omega, 22, 28

H

HAProxy, 41
HashiCorp, 28
HCL, 29
horizontal scale-out, 15

- host, 7
- host mode, 10
- hybrid cloud, 2

I

- IP-per-container, 3
- IPAM, 13, 18
- IPv6, 19
- Istio, 66, 67

K

- Kubernetes
 - architecture, 52
 - East-West traffic, 56
 - Egress, 65
 - Ingress, 62
 - inter-pod networking, 56
 - intra-pod networking, 55
 - network traffic types, 54
 - networking overview, 53
 - North-South traffic, 63
 - service discovery, 59
 - service discovery via DNS, 61
 - service discovery via environment variables, 60
 - service mesh, 66

L

- libnetwork, 46
- Linkerd, 67
- Linux
 - AWS VPC, 16
 - IP per container, 18
 - IPTables, 3
 - IPVLAN, 3
 - namespaces, 3
 - network namespace (reused), 12
 - network namespace (shared), 10
- Linux ELF format, 26
- Linux kernel, 18
- load balancing, 41

M

- Marathon, 26
- Mesos, 26
- Mesos-DNS, 27
 - DNS SRV records, 28

- multi-host container networking
 - overview, 15
- multi-host networking
 - Calico, 16
 - flannel, 16
 - IP address management (IPAM), 18
 - IPv6 (sigh), 19
 - IPVLAN, 18
 - Open vSwitch, 17
 - OpenVPN, 17
 - overlay, 17
 - Weave, 16

N

- networking
 - single host, 7
- NGINX load balancer, 42
- no networking mode, 12
- Nomad, 28

O

- Open Container Initiative (OCI), 7
- OpenStack, 17
- operator, 2
- orchestration, 22

P

- parallel processing, 9
- pets vs cattle, 2
- port allocation, 10, 13
- Project Atomic, 16

R

- Raspberry Pi, 4
- resources
 - CPU, RAM, 23

S

- scheduler, 22
 - constraints, 24
 - quality of service (QoS), 23
- security
 - denial of service attack, 14
- security implications host mode, 11
- service discovery, 22, 33

service mesh
 Conduit, 67
 Istio, 67
SmartStack, 40
Software-Defined Networking, 3
SRV records, 62
Swarm, 25

T

traefik, 42

V

virtual Ethernet bridge, 9
virtual machine, 1
VM, 7
VXLAN, 16

W

Weave, 40