

# CONTENTS IN DETAIL

<b>ACKNOWLEDGMENTS</b>	<b>xvii</b>
------------------------	-------------

<b>INTRODUCTION</b>	<b>xix</b>
---------------------	------------

The Approach .....	xix
Running Examples .....	xx
What You Will Need .....	xx
Run in the Cloud or Local .....	xxi
Terminal Windows .....	xxi

## **PART I MAKING AND USING CONTAINERS**

<b>1</b>	
<b>WHY CONTAINERS MATTER</b>	<b>3</b>

Modern Application Architecture .....	4
Attribute: Cloud Native .....	4
Attribute: Modular .....	4
Attribute: Microservice-Based .....	5
Benefit: Scalability .....	6
Benefit: Reliability .....	6
Benefit: Resilience .....	7
Why Containers .....	7
Requirements for Containers .....	8
Requirements for Orchestration .....	8
Running Containers .....	9
What Containers Look Like .....	9
What Containers Really Are .....	12
Deploying Containers to Kubernetes .....	14
Talking to the Kubernetes Cluster .....	14
Application Overview .....	15
Kubernetes Features .....	16
Final Thoughts .....	17

<b>2</b>		
<b>PROCESS ISOLATION</b>		<b>19</b>
Understanding Isolation .....		20
Why Processes Need Isolation .....		20
File Permissions and Change Root .....		20
Container Isolation .....		22
Container Platforms and Container Runtimes .....		23
Installing containerd .....		23
Using containerd .....		24
Introducing Linux Namespaces .....		25
Containers and Namespaces in CRI-O .....		26
Running Processes in Namespaces Directly .....		30
Final Thoughts .....		32
<b>3</b>		
<b>RESOURCE LIMITING</b>		<b>33</b>
CPU Priorities .....		34
Real-Time and Non-Real-Time Policies .....		34
Setting Process Priorities .....		35
Linux Control Groups .....		37
CPU Quotas with cgroups .....		39
CPU Quota with CRI-O and crictl .....		41
Memory Limits .....		43
Network Bandwidth Limits .....		47
Final Thoughts .....		49
<b>4</b>		
<b>NETWORK NAMESPACES</b>		<b>51</b>
Network Isolation .....		52
Network Namespaces .....		55
Inspecting Network Namespaces .....		55
Creating Network Namespaces .....		57
Bridge Interfaces .....		60
Adding Interfaces to a Bridge .....		61
Tracing Traffic .....		62
Masquerade .....		64
Final Thoughts .....		67

<b>5</b>	<b>CONTAINER IMAGES AND RUNTIME LAYERS</b>	<b>69</b>
Filesystem Isolation .....		70
Container Image Contents .....		70
Image Versions and Layers .....		72
Building Container Images .....		74
Using a Dockerfile .....		74
Tagging and Publishing Images .....		76
Image and Container Storage .....		78
Overlay Filesystems .....		78
Understanding Container Layers .....		80
Practical Image Building Advice .....		81
Open Container Initiative .....		82
Final Thoughts .....		84

## **PART II CONTAINERS IN KUBERNETES**

<b>6</b>	<b>WHY KUBERNETES MATTERS</b>	<b>87</b>
Running Containers in a Cluster .....		88
Cross-Cutting Concerns .....		88
Kubernetes Concepts .....		89
Cluster Deployment .....		90
Prerequisite Packages .....		91
Kubernetes Packages .....		92
Cluster Initialization .....		94
Joining Nodes to the Cluster .....		97
Installing Cluster Add-ons .....		100
Network Driver .....		100
Installing Storage .....		102
Ingress Controller .....		103
Metrics Server .....		104
Exploring a Cluster .....		105
Final Thoughts .....		109

## **7**

### **DEPLOYING CONTAINERS TO KUBERNETES** **111**

Pods .....	111
Deploying a Pod .....	112
Pod Details and Logging .....	114
Deployments .....	115
Creating a Deployment .....	115
Monitoring and Scaling .....	119
Autoscaling .....	120
Other Controllers .....	122
Jobs and CronJobs .....	122
StatefulSets .....	124
Daemon Sets .....	127
Final Thoughts .....	128

## **8**

### **OVERLAY NETWORKS** **129**

Cluster Networking .....	130
CNI Plug-ins .....	130
Pod Networking .....	131
Cross-Node Networking .....	134
Calico Networking .....	135
WeaveNet .....	139
Choosing a Network Plug-in .....	144
Network Customization .....	144
Final Thoughts .....	149

## **9**

### **SERVICE AND INGRESS NETWORKS** **151**

Services .....	152
Creating a Service .....	153
Service DNS .....	155
Name Resolution and Namespaces .....	156
Traffic Routing .....	158
External Networking .....	160
External Services .....	161
Ingress Services .....	162
Ingress in Production .....	165
Final Thoughts .....	166

<b>10</b>		
<b>WHEN THINGS GO WRONG</b>		<b>167</b>
Scheduling .....		167
No Available Nodes .....		168
Insufficient Resources .....		170
Pulling Images .....		173
Running Containers .....		176
Debugging Using Logs .....		177
Debugging Using Exec .....		179
Debugging Using Port Forwarding .....		183
Final Thoughts .....		185
<b>11</b>		
<b>CONTROL PLANE AND ACCESS CONTROL</b>		<b>187</b>
API Server .....		188
API Server Authentication .....		190
Client Certificates .....		190
Bootstrap Tokens .....		192
Service Accounts .....		194
Role-Based Access Controls .....		195
Roles and Cluster Roles .....		195
Role Bindings and Cluster Role Bindings .....		197
Assigning a Service Account to Pods .....		198
Binding Roles to Users .....		201
Final Thoughts .....		203
<b>12</b>		
<b>CONTAINER RUNTIME</b>		<b>205</b>
Node Service .....		205
Kubelet Cluster Configuration .....		207
Kubelet Container Runtime Configuration .....		208
Kubelet Network Configuration .....		209
Static Pods .....		210
Node Maintenance .....		212
Node Draining and Cordonning .....		212
Unhealthy Nodes .....		214
Node Unreachable .....		216
Final Thoughts .....		218

<b>13</b>		
<b>HEALTH PROBES</b>		<b>219</b>
About Probes .....		220
Liveness Probes .....		220
Exec Probes .....		220
HTTP Probes .....		223
TCP Probes .....		225
Startup Probes .....		226
Readiness Probes .....		228
Final Thoughts .....		232

<b>14</b>		
<b>LIMITS AND QUOTAS</b>		<b>233</b>
Requests and Limits .....		234
Processing and Memory Limits .....		234
Cgroup Enforcement .....		236
Network Limits .....		238
Quotas .....		243
Final Thoughts .....		248

<b>15</b>		
<b>PERSISTENT STORAGE</b>		<b>249</b>
Storage Classes .....		249
Storage Class Definition .....		250
CSI Plug-in Internals .....		251
Persistent Volumes .....		252
Stateful Sets .....		252
Volumes and Claims .....		256
Deployments .....		259
Access Modes .....		262
Final Thoughts .....		264

<b>16</b>		
<b>CONFIGURATION AND SECRETS</b>		<b>265</b>
Injecting Configuration .....		265
Externalizing Configuration .....		267
Protecting Secrets .....		269
Injecting Files .....		272
Cluster Configuration Repository .....		274
Using etcdctl .....		275
Deciphering Data in etcd .....		276
Final Thoughts .....		278

**17**  
**CUSTOM RESOURCES AND OPERATORS** **279**

Custom Resources .....	280
Creating CRDs.....	281
Watching CRDs.....	284
Operators .....	288
Final Thoughts .....	292

**PART III**  
**PERFORMANT KUBERNETES**

**18**  
**AFFINITY AND DEVICES** **295**

Affinity and Anti-affinity .....	296
Anti-affinity .....	297
Affinity.....	299
Service Traffic Routing .....	301
Hardware Resources .....	303
Final Thoughts .....	306

**19**  
**TUNING QUALITY OF SERVICE** **309**

Achieving Predictability .....	310
Quality of Service Classes .....	310
BestEffort .....	311
Burstable .....	312
Guaranteed .....	314
QoS Class Eviction .....	315
Choosing a QoS Class .....	316
Pod Priority .....	317
Final Thoughts .....	321

**20**  
**APPLICATION RESILIENCY** **323**

Example Application Stack .....	324
Database .....	324
Application Deployment .....	326
Pod Autoscaling .....	329
Application Service .....	330

Application and Cluster Monitoring .....	332
Prometheus Monitoring .....	333
Deploying kube-prometheus .....	334
Cluster Metrics .....	337
Adding Monitoring for Services .....	339
Final Thoughts .....	343

**INDEX** **345**