
Table of Contents

Preface	ix
1. Introduction	1
Setting the Scene	2
Starting to Threat Model	3
Threat Actors	4
Your First Threat Model	7
Attack Trees	9
Example Attack Trees	11
Prior Art	13
Conclusion	13
2. Pod-Level Resources	15
Defaults	15
Threat Model	16
Anatomy of the Attack	17
Remote Code Execution	18
Network Attack Surface	19
Kubernetes Workloads: Apps in a Pod	20
What's a Pod?	22
Understanding Containers	27
Sharing Network and Storage	28
What's the Worst That Could Happen?	30
Container Breakout	34
Pod Configuration and Threats	37
Pod Header	37
Reverse Uptime	38
Labels	39

Managed Fields	39
Pod Namespace and Owner	40
Environment Variables	40
Container Images	41
Pod Probes	43
CPU and Memory Limits and Requests	43
DNS	44
Pod securityContext	46
Pod Service Accounts	49
Scheduler and Tolerations	49
Pod Volume Definitions	49
Pod Network Status	50
Using the securityContext Correctly	50
Enhancing the securityContext with Kubesec	52
Hardened securityContext	53
Into the Eye of the Storm	57
Conclusion	58
3. Container Runtime Isolation.....	59
Defaults	59
Threat Model	60
Containers, Virtual Machines, and Sandboxes	62
How Virtual Machines Work	64
Benefits of Virtualization	67
What's Wrong with Containers?	67
User Namespace Vulnerabilities	69
Sandboxing	73
gVisor	75
Firecracker	82
Kata Containers	84
rust-vmm	85
Risks of Sandboxing	86
Kubernetes Runtime Class	87
Conclusion	88
4. Applications and Supply Chain.....	89
Defaults	90
Threat Model	90
The Supply Chain	91
Software	94
Scanning for CVEs	95
Ingesting Open Source Software	96

Which Producers Do We Trust?	97
CNCF Security Technical Advisory Group	98
Architecting Containerized Apps for Resilience	98
Detecting Trojans	99
Captain Hashjack Attacks a Supply Chain	100
Post-Compromise Persistence	102
Risks to Your Systems	102
Container Image Build Supply Chains	103
Software Factories	103
Blessed Image Factory	104
Base Images	105
The State of Your Container Supply Chains	106
Third-Party Code Risk	107
Software Bills of Materials	108
Human Identity and GPG	110
Signing Builds and Metadata	110
Notary v1	111
sigstore	111
in-toto and TUF	113
GCP Binary Authorization	113
Grafeas	114
Infrastructure Supply Chain	114
Operator Privileges	114
Attacking Higher Up the Supply Chain	114
Types of Supply Chain Attack	115
Open Source Ingestion	117
Application Vulnerability Throughout the SDLC	119
Defending Against SUNBURST	120
Conclusion	123
5. Networking.....	125
Defaults	126
Intra-Pod Networking	128
Inter-Pod Traffic	128
Pod-to-Worker Node Traffic	129
Cluster-External Traffic	129
The State of the ARP	130
No securityContext	131
No Workload Identity	132
No Encryption on the Wire	132
Threat Model	133
Traffic Flow Control	134

The Setup	134
Network Policies to the Rescue!	137
Service Meshes	139
Concept	139
Options and Uptake	140
Case Study: mTLS with Linkerd	141
eBPF	144
Concept	144
Options and Uptake	144
Case Study: Attaching a Probe to a Go Program	145
Conclusion	147
6. Storage.....	149
Defaults	150
Threat Model	150
Volumes and Datastores	152
Everything Is a Stream of Bytes	152
What's a Filesystem?	153
Container Volumes and Mounts	154
OverlayFS	155
tmpfs	156
Volume Mount Breaks Container Isolation	158
The /proc/self/exe CVE	160
Sensitive Information at Rest	162
Mounted Secrets	162
Attacking Mounted Secrets	163
Storage Concepts	164
Container Storage Interface	164
Projected Volumes	165
Attacking Volumes	167
The Dangers of Host Mounts	169
Other Secrets and Exfiltrating from Datastores	169
Conclusion	170
7. Hard Multitenancy.....	171
Defaults	172
Threat Model	172
Namespaced Resources	173
Node Pools	174
Node Taints	176
Soft Multitenancy	177
Hard Multitenancy	178

Hostile Tenants	178
Sandboxing and Policy	179
Public Cloud Multitenancy	180
Control Plane	181
API Server and etcd	182
Scheduler and Controller Manager	184
Data Plane	187
Cluster Isolation Architecture	188
Cluster Support Services and Tooling Environments	190
Security Monitoring and Visibility	191
Conclusion	191
8. Policy.....	193
Types of Policies	194
Defaults	194
Network Traffic	195
Limiting Resource Allocations	195
Resource Quotas	196
Runtime Policies	197
Access Control Policies	197
Threat Model	198
Common Expectations	198
Breakglass Scenario	199
Auditing	199
Authentication and Authorization	200
Human Users	201
Workload Identity	201
Role-Based Access Control (RBAC)	204
RBAC Recap	204
A Simple RBAC Example	205
Authoring RBAC	207
Analyzing and Visualizing RBAC	209
RBAC-Related Attacks	211
Generic Policy Engines	212
Open Policy Agent	212
Kyverno	218
Other Policy Offerings	220
Conclusion	221
9. Intrusion Detection.....	223
Defaults	223
Threat Model	224

Traditional IDS	224
eBPF-Based IDS	226
Kubernetes and Container Intrusion Detection	227
Falco	227
Machine Learning Approaches to IDS	229
Container Forensics	230
Honeypots	232
Auditing	234
Detection Evasion	235
Security Operations Centers	236
Conclusion	237
10. Organizations.....	239
The Weakest Link	240
Cloud Providers	241
Shared Responsibility	242
Account Hygiene	243
Grouping People and Resources	245
Other Considerations	246
On-Premises Environments	247
Common Considerations	249
Threat Model Explosion	249
How SLOs Can Put Additional Pressure on You	252
Social Engineering	252
Privacy and Regulatory Concerns	254
Conclusion	255
A. A Pod-Level Attack.....	257
B. Resources.....	271
Index.....	279