

Contents

Copyright	6
About the author	7
Acknowledgements	8
Preface	9
0.1 Who should read this book	10
1 Introduction	11
1.1 Communication	12
1.2 Coordination	13
1.3 Scalability	13
1.4 Resiliency	15
1.5 Operations	16
1.6 Anatomy of a distributed system	17
 I Communication	 20
2 Reliable links	23
2.1 Reliability	23
2.2 Connection lifecycle	24
2.3 Flow control	25
2.4 Congestion control	27
2.5 Custom protocols	28
3 Secure links	30
3.1 Encryption	30
3.2 Authentication	31
3.3 Integrity	33

3.4	Handshake	34
4	Discovery	35
5	APIs	39
5.1	HTTP	41
5.2	Resources	43
5.3	Request methods	45
5.4	Response status codes	46
5.5	OpenAPI	48
5.6	Evolution	49
II	Coordination	51
6	System models	54
7	Failure detection	57
8	Time	59
8.1	Physical clocks	59
8.2	Logical clocks	61
8.3	Vector clocks	63
9	Leader election	66
9.1	Raft leader election	66
9.2	Practical considerations	68
10	Replication	72
10.1	State machine replication	73
10.2	Consensus	76
10.3	Consistency models	77
10.3.1	Strong consistency	78
10.3.2	Sequential consistency	80
10.3.3	Eventual consistency	80
10.3.4	CAP theorem	82
10.4	Chain replication	83
10.5	Solving the CAP theorem	87
10.5.1	Broadcast protocols	88
10.5.2	Conflict free replicated data types	90
10.5.3	Dynamo-style data stores	94
10.5.4	CALM theorem	97
10.5.5	Causal consistency	99
10.6	Coordination avoidance	101

11 Transactions	103
11.1 ACID	103
11.2 Isolation	105
11.2.1 Concurrency control	107
11.3 Atomicity	108
11.3.1 Two-phase commit	109
11.4 Asynchronous transactions	111
11.4.1 Log-based transactions	112
11.4.2 Sagas	115
11.4.3 Isolation	118
 III Scalability	 119
12 Functional decomposition	123
12.1 Microservices	123
12.1.1 Benefits	125
12.1.2 Costs	126
12.1.3 Practical considerations	128
12.2 API gateway	130
12.2.1 Routing	130
12.2.2 Composition	131
12.2.3 Translation	132
12.2.4 Cross-cutting concerns	132
12.2.5 Caveats	135
12.3 CQRS	137
12.4 Messaging	138
12.4.1 Guarantees	141
12.4.2 Exactly-once processing	143
12.4.3 Failures	144
12.4.4 Backlogs	144
12.4.5 Fault isolation	145
12.4.6 Reference plus blob	146
 13 Partitioning	 148
13.1 Sharding strategies	148
13.1.1 Range partitioning	149
13.1.2 Hash partitioning	149
13.2 Rebalancing	153
13.2.1 Static partitioning	153
13.2.2 Dynamic partitioning	153
13.2.3 Practical considerations	154

14 Duplication	155
14.1 Network load balancing	155
14.1.1 DNS load balancing	158
14.1.2 Transport layer load balancing	158
14.1.3 Application layer load balancing	161
14.1.4 Geo load balancing	163
14.2 Replication	165
14.2.1 Single leader replication	165
14.2.2 Multi-leader replication	168
14.2.3 Leaderless replication	170
14.3 Caching	171
14.3.1 Policies	171
14.3.2 In-process cache	172
14.3.3 Out-of-process cache	173
 IV Resiliency	 176
15 Common failure causes	179
15.1 Single point of failure	179
15.2 Unreliable network	180
15.3 Slow processes	181
15.4 Unexpected load	182
15.5 Cascading failures	183
15.6 Risk management	184
 16 Downstream resiliency	 186
16.1 Timeout	186
16.2 Retry	189
16.2.1 Exponential backoff	190
16.2.2 Retry amplification	191
16.3 Circuit breaker	192
16.3.1 State machine	193
 17 Upstream resiliency	 195
17.1 Load shedding	195
17.2 Load leveling	196
17.3 Rate-limiting	197
17.3.1 Single-process implementation	199
17.3.2 Distributed implementation	202
17.4 Bulkhead	203
17.5 Health endpoint	206
17.5.1 Health checks	206

17.6 Watchdog	207
V Testing and operations	209
18 Testing	212
18.1 Scope	213
18.2 Size	215
18.3 Practical considerations	217
19 Continuous delivery and deployment	219
19.1 Review and build	220
19.2 Pre-production	222
19.3 Production	222
19.4 Rollbacks	223
20 Monitoring	226
20.1 Metrics	227
20.2 Service-level indicators	230
20.3 Service-level objectives	233
20.4 Alerts	235
20.5 Dashboards	238
20.5.1 Best practices	240
20.6 On-call	241
21 Observability	243
21.1 Logs	244
21.2 Traces	247
21.3 Putting it all together	249
22 Final words	251