

Contents

Copyright	ix
About the author	xi
Acknowledgements	xiii
Preface	xv
1 Introduction	1
1.1 Communication	2
1.2 Coordination	3
1.3 Scalability	3
1.4 Resiliency	5
1.5 Maintainability	6
1.6 Anatomy of a distributed system	7
I Communication	11
2 Reliable links	17
2.1 Reliability	18
2.2 Connection lifecycle	18
2.3 Flow control	20
2.4 Congestion control	21
2.5 Custom protocols	23
3 Secure links	25
3.1 Encryption	25
3.2 Authentication	26
3.3 Integrity	28
3.4 Handshake	29

4	Discovery	31
5	APIs	35
5.1	HTTP	37
5.2	Resources	39
5.3	Request methods	41
5.4	Response status codes	42
5.5	OpenAPI	43
5.6	Evolution	45
5.7	Idempotency	46
II	Coordination	53
6	System models	57
7	Failure detection	61
8	Time	63
8.1	Physical clocks	63
8.2	Logical clocks	65
8.3	Vector clocks	67
9	Leader election	71
9.1	Raft leader election	72
9.2	Practical considerations	73
10	Replication	77
10.1	State machine replication	78
10.2	Consensus	81
10.3	Consistency models	83
10.4	Chain replication	90
11	Coordination avoidance	95
11.1	Broadcast protocols	96
11.2	Conflict-free replicated data types	98
11.3	Dynamo-style data stores	103
11.4	The CALM theorem	105
11.5	Causal consistency	106
11.6	Practical considerations	110
12	Transactions	111
12.1	ACID	112
12.2	Isolation	113

12.3	Atomicity	119
12.4	NewSQL	122
13	Asynchronous transactions	127
13.1	Outbox pattern	128
13.2	Sagas	130
13.3	Isolation	133
III	Scalability	137
14	HTTP caching	145
14.1	Reverse proxies	148
15	Content delivery networks	151
15.1	Overlay network	151
15.2	Caching	153
16	Partitioning	155
16.1	Range partitioning	157
16.2	Hash partitioning	158
17	File storage	163
17.1	Blob storage architecture	163
18	Network load balancing	169
18.1	DNS load balancing	174
18.2	Transport layer load balancing	175
18.3	Application layer load balancing	178
19	Data storage	181
19.1	Replication	181
19.2	Partitioning	184
19.3	NoSQL	185
20	Caching	191
20.1	Policies	192
20.2	Local cache	193
20.3	External cache	194
21	Microservices	197
21.1	Caveats	199
21.2	API gateway	202

22	Control planes and data planes	209
22.1	Scale imbalance	211
22.2	Control theory	214
23	Messaging	217
23.1	Guarantees	221
23.2	Exactly-once processing	223
23.3	Failures	224
23.4	Backlogs	224
23.5	Fault isolation	225
IV	Resiliency	229
24	Common failure causes	233
24.1	Hardware faults	233
24.2	Incorrect error handling	234
24.3	Configuration changes	234
24.4	Single points of failure	235
24.5	Network faults	236
24.6	Resource leaks	237
24.7	Load pressure	238
24.8	Cascading failures	238
24.9	Managing risk	240
25	Redundancy	243
25.1	Correlation	244
26	Fault isolation	247
26.1	Shuffle sharding	248
26.2	Cellular architecture	250
27	Downstream resiliency	253
27.1	Timeout	253
27.2	Retry	255
27.3	Circuit breaker	258
28	Upstream resiliency	261
28.1	Load shedding	261
28.2	Load leveling	262
28.3	Rate-limiting	263
28.4	Constant work	269

V	Maintainability	275
29	Testing	279
29.1	Scope	280
29.2	Size	281
29.3	Practical considerations	283
29.4	Formal verification	285
30	Continuous delivery and deployment	289
30.1	Review and build	290
30.2	Pre-production	292
30.3	Production	293
30.4	Rollbacks	293
31	Monitoring	297
31.1	Metrics	298
31.2	Service-level indicators	301
31.3	Service-level objectives	304
31.4	Alerts	306
31.5	Dashboards	308
31.6	Being on call	312
32	Observability	315
32.1	Logs	316
32.2	Traces	319
32.3	Putting it all together	321
33	Manageability	323
34	Final words	327