

# Kafka: The Definitive Guide

*Real-Time Data and Stream Processing at Scale*

*Neha Narkhede, Gwert Shapira, and Todd Palino*

# Table of Contents

Foreword	xi
Preface	xv
1. Meet Kafka	1
Publish/Subscribe Messaging	1
How It Starts	2
Individual Queue Systems	3
Enter Kafka	4
Messages and Batches	4
Schemas	5
Topics and Partitions	5
Producers and Consumers	6
Brokers and Clusters	7
Multiple Clusters	8
Why Kafka?	10
Multiple Producers	10
Multiple Consumers	10
Disk-Based Retention	10
Scalable	10
High Performance	11
The Data Ecosystem	11
Use Cases	12
Kafka's Origin	14
LinkedIn's Problem	14
The Birth of Kafka	15
Open Source	15
The Name	16

Getting Started with Kafka	16
2. Installing Kafka	17
First Things First	17
Choosing an Operating System	17
Installing Java	17
Installing Zookeeper	18
Installing a Kafka Broker	20
Broker Configuration	21
General Broker	21
Topic Defaults	24
Hardware Selection	28
Disk Throughput	29
Disk Capacity	29
Memory	29
Networking	30
CPU	30
Kafka in the Cloud	30
Kafka Clusters	31
How Many Brokers?	32
Broker Configuration	32
OS Tuning	32
Production Concerns	36
Garbage Collector Options	36
Datacenter Layout	37
Colocating Applications on Zookeeper	37
Summary	39
3. Kafka Producers: Writing Messages to Kafka	41
Producer Overview	42
Constructing a Kafka Producer	44
Sending a Message to Kafka	46
Sending a Message Synchronously	46
Sending a Message Asynchronously	47
Configuring Producers	48
Serializers	52
Custom Serializers	52
Serializing Using Apache Avro	54
Using Avro Records with Kafka	56
Partitions	59
Old Producer APIs	61
Summary	62

4. Kafka Consumers: Reading Data from Kafka	63
Kafka Consumer Concepts	63
Consumers and Consumer Groups	63
Consumer Groups and Partition Rebalance	66
Creating a Kafka Consumer	68
Subscribing to Topics	69
The Poll Loop	70
Configuring Consumers	72
Commits and Offsets	75
Automatic Commit	76
Commit Current Offset	77
Asynchronous Commit	78
Combining Synchronous and Asynchronous Commits	80
Commit Specified Offset	80
Rebalance Listeners	82
Consuming Records with Specific Offsets	84
But How Do We Exit?	86
Deserializers	88
Standalone Consumer: Why and How to Use a Consumer Without a Group	92
Older Consumer APIs	93
Summary	93
5. Kafka Internals	95
Cluster Membership	95
The Controller	96
Replication	97
Request Processing	99
Produce Requests	101
Fetch Requests	102
Other Requests	104
Physical Storage	105
Partition Allocation	106
File Management	107
File Format	108
Indexes	109
Compaction	110
How Compaction Works	110
Deleted Events	112
When Are Topics Compacted?	112
Summary	113

6. Reliable Data Delivery	115
Reliability Guarantees	116
Replication	117
Broker Configuration	118
Replication Factor	118
Unclean Leader Election	119
Minimum In-Sync Replicas	121
Using Producers in a Reliable System	121
Send Acknowledgments	122
Configuring Producer Retries	123
Additional Error Handling	124
Using Consumers in a Reliable System	125
Important Consumer Configuration Properties for Reliable Processing	126
Explicitly Committing Offsets in Consumers	127
Validating System Reliability	129
Validating Configuration	130
Validating Applications	131
Monitoring Reliability in Production	131
Summary	133
7. Building Data Pipelines	135
Considerations When Building Data Pipelines	136
Timeliness	136
Reliability	137
High and Varying Throughput	137
Data Formats	138
Transformations	139
Security	139
Failure Handling	140
Coupling and Agility	140
When to Use Kafka Connect Versus Producer and Consumer	141
Kafka Connect	142
Running Connect	142
Connector Example: File Source and File Sink	144
Connector Example: MySQL to Elasticsearch	146
A Deeper Look at Connect	151
Alternatives to Kafka Connect	154
Ingest Frameworks for Other Datastores	155
GUI-Based ETL Tools	155
Stream-Processing Frameworks	155
Summary	156

8. Cross-Cluster Data Mirroring	157
Use Cases of Cross-Cluster Mirroring	158
Multicluster Architectures	158
Some Realities of Cross-Datacenter Communication	159
Hub-and-Spokes Architecture	160
Active-Active Architecture	161
Active-Standby Architecture	163
Stretch Clusters	169
Apache Kafka's MirrorMaker	170
How to Configure	171
Deploying MirrorMaker in Production	172
Tuning MirrorMaker	175
Other Cross-Cluster Mirroring Solutions	178
Uber uReplicator	178
Confluents Replicator	179
Summary	180
9. Administering Kafka	181
Topic Operations	181
Creating a New Topic	182
Adding Partitions	183
Deleting a Topic	184
Listing All Topics in a Cluster	185
Describing Topic Details	185
Consumer Groups	186
List and Describe Groups	186
Delete Group	188
Offset Management	188
Dynamic Configuration Changes	190
Overriding Topic Configuration Defaults	190
Overriding Client Configuration Defaults	192
Describing Configuration Overrides	192
Removing Configuration Overrides	193
Partition Management	193
Preferred Replica Election	193
Changing a Partition's Replicas	195
Changing Replication Factor	198
Dumping Log Segments	199
Replica Verification	201
Consuming and Producing	202
Console Consumer	202
Console Producer	205

Client ACLs	207
Unsafe Operations	207
Moving the Cluster Controller	208
Killing a Partition Move	208
Removing Topics to Be Deleted	209
Deleting Topics Manually	209
Summary	210
10. Monitoring Kafka	211
Metric Basics	211
Where Are the Metrics?	211
Internal or External Measurements	212
Application Health Checks	213
Metric Coverage	213
Kafka Broker Metrics	213
Under-Replicated Partitions	214
Broker Metrics	220
Topic and Partition Metrics	229
JVM Monitoring	231
OS Monitoring	232
Logging	235
Client Monitoring	236
Producer Metrics	236
Consumer Metrics	239
Quotas	242
Lag Monitoring	243
End-to-End Monitoring	244
Summary	244
11. Stream Processing	247
What Is Stream Processing?	248
Stream-Processing Concepts	251
Time	251
State	252
Stream-Table Duality	253
Time Windows	255
Stream-Processing Design Patterns	256
Single-Event Processing	256
Processing with Local State	257
Multiphase Processing/Repartitioning	259
Processing with External Lookup: Stream-Table Join	260
Streaming Join	261

Out-of-Sequence Events	262
Reprocessing	264
Kafka Streams by Example	264
Word Count	265
Stock Market Statistics	268
Click Stream Enrichment	270
Kafka Streams: Architecture Overview	272
Building a Topology	272
Scaling the Topology	273
Surviving Failures	276
Stream Processing Use Cases	277
How to Choose a Stream-Processing Framework	278
Summary	280
A. Installing Kafka on Other Operating Systems	281
Index	287