

# Kafka Cluster Minimum Requirements

Last Modified on 03/01/2020 10:09 am EST

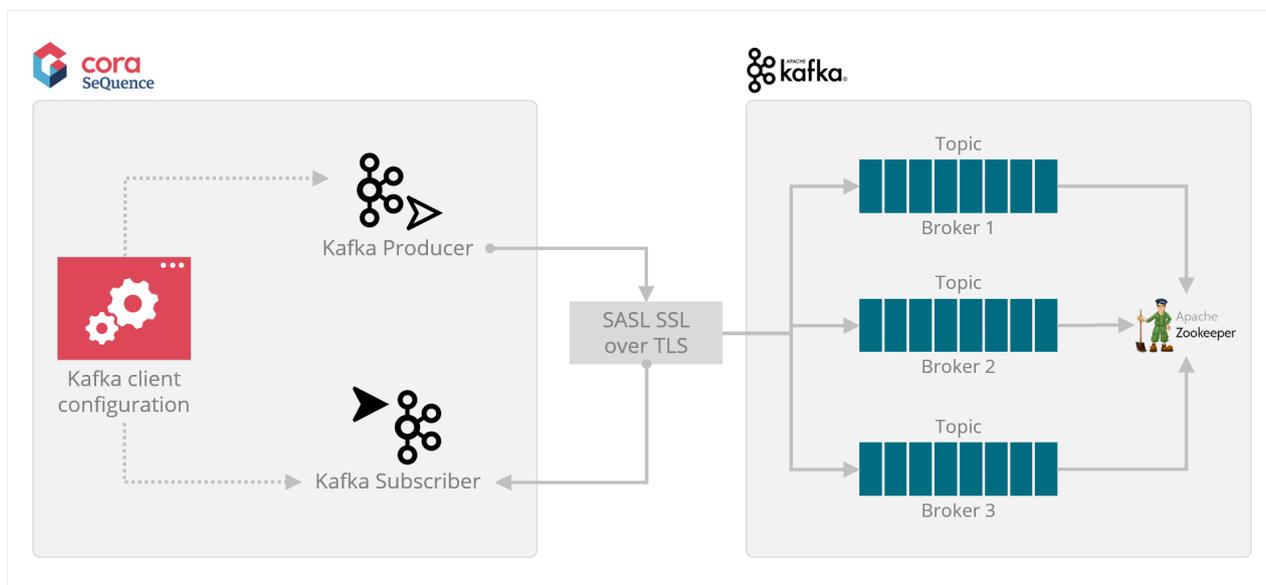
## V9.2.2

### Overview

You can use Apache Kafka® to send messages within Cora SeSequence or between Cora SeSequence and other applications. Apache Kafka is a distributed streaming platform that:

- Publishes and subscribes to streams of records, similar to a message queue or enterprise messaging system.
- Stores streams of records in a fault-tolerant durable way.
- Processes streams of records as they occur.

To integrate with Cora SeSequence and maintain a robust, fault-tolerant and high-availability Kafka cluster, follow the recommendations below.



Component	Recommendation	Details
Kafka brokers	Three brokers	To maintain high-availability of the Kafka cluster (broker leader re-election), and to provide partition replication and better partitioning of topics.
Apache ZooKeeper™	One node	Zookeeper coordinates between the brokers (controller election), keeps topic configurations, and access control lists (ACL).
Security protocol	SASL_SSL protocol	To enable Kafka authentication and encryption.

### Terminology

Make sure that you are aware of the terminology used in Cora SeQuence.

<b>Cora SeQuence activity</b>	<b>Kafka component</b>
Kafka Producer	Producer
Kafka Subscriber	Consumer

For more details on Kafka components, see [this page](#).

### Minimum system requirements to run Kafka in a production environment

<b>Item</b>	<b>Requirement</b>	<b>Details</b>
Memory	8 GB RAM	Kafka relies heavily on the file system for storing and caching messages. Kafka uses heap space very carefully and does not require setting heap sizes more than 6 GB.
CPU	4 cores	Note that if SSL is enabled, the CPU requirements can be significantly higher. The exact requirements depend on the CPU type and JVM implementation.  If you need to choose between faster CPUs or more cores, choose more cores. The extra concurrency that multiple cores offers will far outweigh a slightly faster clock speed.
Disk storage size	500 GB	Use multiple drives to maximize throughput. Avoid network-attached storage (NAS).
Network	1 GbE-10 GbE	Avoid clusters that span multiple data centers.

### Minimum software version

<b>Software</b>	<b>Version</b>
Apache Kafka	2.3.1-0
Apache Zookeeper	3.4.14
Java	1.8 and later

For more information on Kafka system requirements, see [this page](#).