

Read-Write Separation Deployment for Kylin 4.0

- [Background](#)
- [Architecture](#)
- [Preparation](#)
- [Deploy Steps](#)
- [Conclusion](#)
- [FAQ](#)

Background

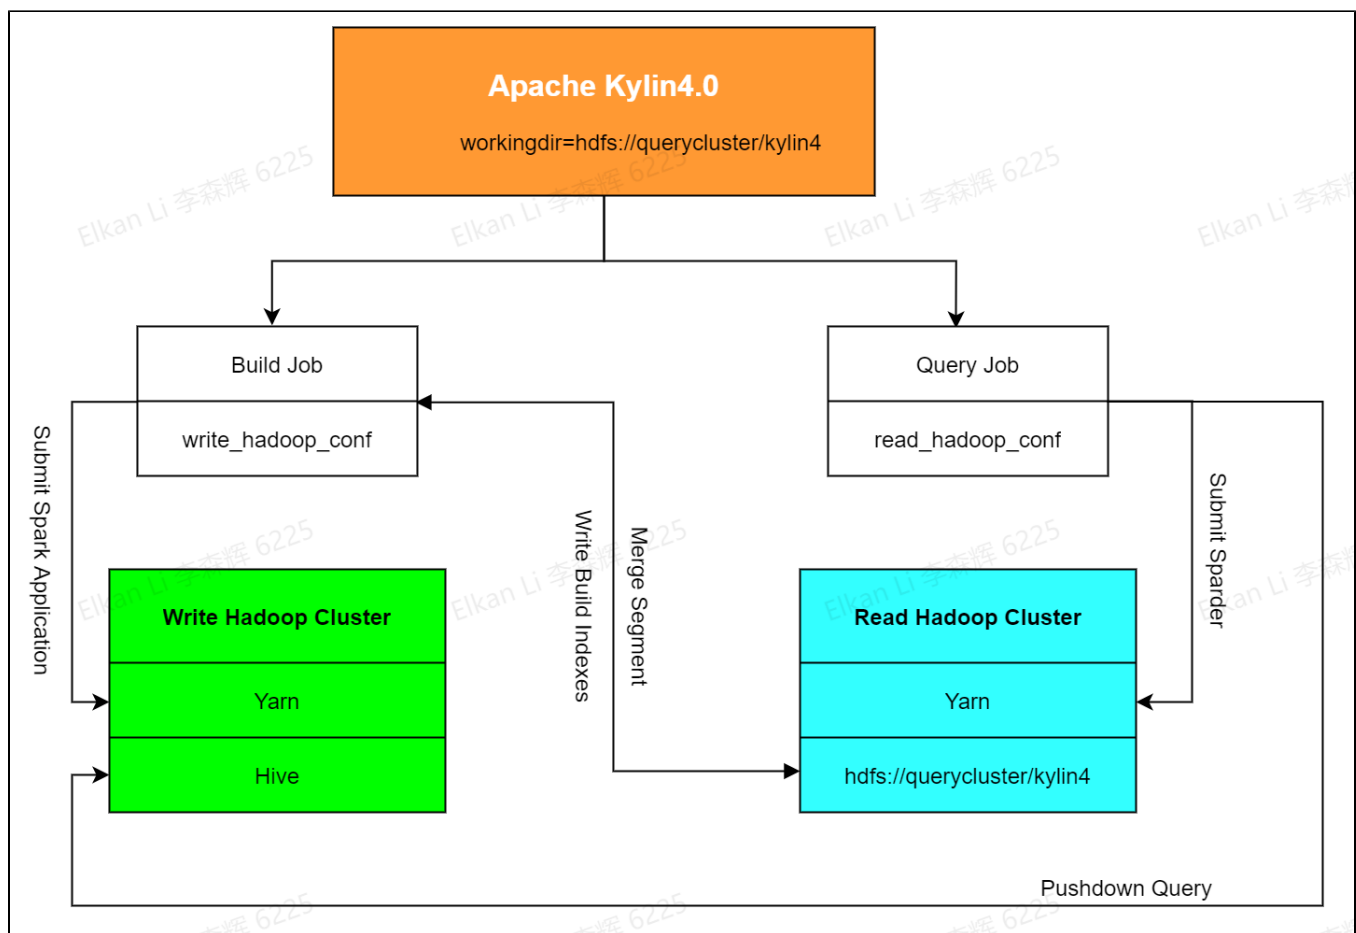
You maybe know that Kylin 4.0 is using Spark engine to build cube and query cube, if the build (spark) jobs and query (spark) jobs are both running in the same Hadoop cluster, the performance of build/query will be affected because of the **resource competition**.

Now, Kylin 4.0 supports running build jobs and query jobs on the two Hadoop clusters. The build (spark) jobs will be sent to **build cluster** to build cube data, and then the cube/cuboid data will be wrote to the HDFS of the **query cluster** directly, so that query related workload will be move to **query cluster**.

With a read-write separation deployment, we can separate build and query computation workload(mostly the yarn resource and disk I/O). In the current implementation, HDFS of **query cluster** is not exclusive to the **query cluster** only, because **build cluster** will read data from **query cluster**'s HDFS) when merging segments.

Architecture

The architecture and workflow for read-write separation in Kylin4.0 will like below:



Note

1. Kylin 4.0 uses the Yarn resources on build cluster to build cube data and then write the cube data back to the HDFS on query cluster directly.
2. Build jobs read the Hive data sources which are on build cluster.
3. When executing pushdown query, Kylin 4.0 reads the Hive data from build cluster.
4. Before merge segments the merge job will read cube data from query cluster.

Preparation

Before you starting deploy that need make sure your hadoop cluster environment should be supported Kylin4.0, also the hadoop version of build cluster and query cluster must be the same. In addition to noted the following environment as well:

1. Verify that can access between build cluster and query cluster resources(mostly the HDFS service).
2. Make sure the commands like hdfs and hive are all working properly and can access cluster resources.
3. If the two clusters have enabled the HDFS NameNode HA, please **make sure their HDFS Nameservice names are different**. If they are the same, please change one of them to avoid conflict.
4. Please make sure the network latency between the two clusters is low enough, as there will be a large number of data moved back and forth during model build process.

Deploy Steps

At first you can install Kylin4.0 follow in the official guide in one node which contains Hadoop client configuration of **query cluster**. Then configure the parameters needed for the build action as below:

1. Create a directory called build_hadoop_conf in the \$KYLIN_HOME and then copy the hadoop configuration files of the **build cluster** into this directory (**Note**: make sure copy the real configuration files, not the symbolic links).
2. Set the value of the configuration 'kylin.engine.submit-hadoop-conf-dir' in the '\$KYLIN_HOME/conf/kylin.properties' to the directory created in the above step.
3. Copy the [hive-site.xml](#) from **build cluster** to **query cluster** into right directory (typically, hive configuration should located at : [/etc/hive/conf](#)). So the query can load Hive metadata from **build cluster** properly.
4. The value of the configuration 'kylin.engine.spark-conf.spark.yarn.queue' in the '\$KYLIN_HOME/conf/kylin.properties' should be configured as the queue of the Yarn on **build cluster**.

Conclusion

Kylin4.0 doesn't use HBase as storage layer which makes the architecture more simple, so that we can easily deploy the read-write separation for Kylin. With this deployment we can get a more stable Kylin cluster which can provide high performance data service ability. It is suitable for scenarios such as the use of the dashboard, self-service analysis, data interface service and so on which need highly query performance.

FAQ

Current in this read-write separation deployment there have some limitations and FAQ list as below:

Q: Which scripts or command can't be supported?

A: \$KYLIN_HOME/bin/sample.sh is not supported in this deployment mode.

Q: Whether the configuration files of the Hadoop cluster are automatically updated synchronously?

A: There needs to be an upgrade to build cluster's configuration files by manual.

Q: Whether need to restart the cluster when configuring Hadoop communication parameters between build cluster and query cluster?

A: Yes, Please be careful in this setting in the production environment.

Q: Does the build cluster had save cube data as repetition?

A: No, the cube data only store in query cluster.