

Kafka Tiered Storage Early Access Release Notes

This is an early access release for tiered storage feature([KIP-405](#)), and it is not recommended for use in production environments. Instead, we advise users to create new clusters with version 3.6.0 and test the feature there. Alternatively, you can try the tiered storage feature on existing non-production clusters running version 2.8.0 or later, upgraded to version 3.6.0.

The following features related to Tiered Storage are now available in the early access release:

- Remote Storage at Cluster and Topic Level:** To leverage the remote tier archival and retrieval capabilities, you need to enable remote storage at both at the cluster level and the specific topic level. By enabling remote storage at the cluster level, you activate the feature for the entire Kafka cluster. Subsequently, by enabling remote storage for a particular topic, you designate it for archival and retrieval through remote storage.
- Seamless Client Compatibility:** With Tiered Storage enabled for a topic, no changes are required in Kafka clients to read from that topic.
- Non-intrusive Integration:** We have ensured that the "early access code" for tiered storage feature does not cause any disruptions to other Kafka functionalities when it is disabled at a cluster level. Your existing Kafka setup will continue to function smoothly, as the default configuration for "remote.log.storage.system.enable" is set to false, ensuring the tiered storage feature remains inactive.
- LocalTieredStorage Implementation:** As part of the default implementation, we introduce 'LocalTieredStorage,' a local file-based RemoteStorageManager. LocalTieredStorage facilitates the simulation of remote storage behavior in a controlled and isolated environment during testing.
- Monitoring for Remote Storage feature:** This feature incorporates several new metrics to help you keep track of tiered storage operations effectively. These metrics are available exclusively when the Remote Storage feature is enabled for your cluster.
- Cluster Upgrades:** Clusters upgrading from version any previous version to version 3.6.0 can enable Tiered Storage for topics created after the upgrade to 3.6.0. Topics created on or after version 2.8.0 are also eligible for Tiered Storage. However, for topics created before version 2.8.0, Tiered Storage cannot be enabled after upgrading the cluster to version 3.6.0. To utilize Tiered Storage for older topics, manual steps are required. Specifically, older segments need to be deleted before Tiered Storage can be activated on these topics. Our code does not automatically block this process, leaving the responsibility in the hands of the user.
- Enabling and Disabling Tiered Storage for a Cluster:** We now offer the ability to enable or disable Tiered Storage for an entire Kafka cluster. However, please be aware that disabling Tiered Storage has some limitations, as described in the limitations section below.
- Data Deletion upon Topic Deletion:** When you delete a topic that is utilizing Tiered Storage, the data associated with that topic will be automatically deleted from remote storage.

While the early access release of Tiered Storage offers the opportunity to try out this new feature, it is important to be aware of the following limitations:

- Clusters with Multiple Log Directories:** Tiered storage is not supported on clusters utilizing the multiple Log directories on a broker (JBOD feature). When attempting to enable remote storage on a broker with multiple Log directories (JBOD feature), you will receive a configuration exception.
- Compacted Topics:** Currently, tiered storage is not available for compacted topics. If you attempt to enable remote storage on a compacted topic, you will receive a configuration exception. Also, if this topic WAS a compacted topic, and later updated as a non-compacted topic. In this case, enabling remote storage will not throw configuration exceptions. However, this is still not supported because we assume the topics are not compacted.
- Disabling Tiered Storage for a Topic:** Once Tiered Storage is enabled for a topic, it becomes a permanent configuration and cannot be disabled. The only way to remove the tiered storage feature from a topic is by deleting the topic itself. It is essential to be aware that deleting the topic will result in the data being removed from remote storage. Therefore, exercise caution when enabling Tiered Storage for topics, as it becomes an irreversible operation. In future versions, [KIP-950](#) will remove this limitation and add flexibility to disable / re-enable remote storage for a topic.
- Disabling Tiered Storage at the Cluster Level:** Disabling Tiered Storage for the entire cluster requires manual deletion of all topics using Tiered Storage. Attempting to disable tiered storage at the cluster level without deleting the topics using tiered storage will result in an exception.
- Client Compatibility:** All Kafka clients, regardless of their version, can continue to produce and consume records from topics utilizing Tiered Storage. However, clients with versions earlier than 3.0 are limited in performing administrative actions, such as enabling Tiered Storage on a topic (ex: they might change directly in ZK using `--zookeeper` option). To successfully enable Tiered Storage for a topic, clients must be running Kafka version 3.0 or later, as administrative actions related to Tiered Storage are only supported on clients from version 3.0 onwards.

For the latest information regarding known issues, their resolutions, and possible workarounds, please visit the parent tracker for next release of Tiered

Storage feature at



Unable to render Jira issues macro, execution error.

. We are committed to addressing these issues and providing

a reliable Tiered Storage experience, and your feedback is incredibly valuable in helping us improve the feature.