

Parallelism Tuning

If the Split Spots on the node are not properly configured users may see jobs lag due to resource conflict. When all of the available Split Slots in the cluster are occupied incoming jobs will be delayed until the resources become free. If the Split Size Allocation is not sufficient to store the dataset then the incoming jobs may fail if the cluster is not able to grow any further.

Number of Splits in Cluster	node-scheduler.max-splits-per-node * Number of Nodes
Initial Split Size Allocation for Dataset	hive.max-initial-splits * hive.max-initial-split-size
Maximum Split Size Allocation for Dataset	hive.max-initial-splits * hive.max-split-size

Memory Tuning

Presto features three Memory Pools to manage the available resources and two of the Memory Pools are responsible for data processing. All queries are initially submitted to the General Memory Pool however if the memory limit is exceeded then queries will get blocked until other queries finish and release the memory. If the General Memory Pool limit is reached Presto will move the biggest query running at the moment to Reserved Pool.

maxHeap	The JVM container size.	Defaults to up to 70% of Instance Memory
System Memory	The overhead allocation.	Defaults to 40% of maxHeap
Reserved Memory	Reserved for Largest Job	query.max-memory-per-node
General Memory	First stop for all Jobs	Max Heap - Reserved Memory
Query Memory	Max memory for the Job	set session query_max_memory
Query Memory Limit	Memory used across Nodes	query.max-memory

Avoiding Issues

Presto will delay jobs when there are not enough Split Slots to support the Dataset and jobs will Fail when there is not enough memory to process the query. If any of the below apply to the current environment then the configuration is not powerful enough and users can expect job lag and failure.

Reserved Memory * Number of Nodes < Peak Job Size	Use larger instance.
General Memory * Number of Nodes < Average Job Size * Concurrent Jobs	Use larger instance.
Reserved Memory * Number of Nodes < Query Memory	Adjust Setting.
Reserved Memory * Number of Nodes < Query Memory Limit	Adjust Setting.