

BatchSize, ChannelCapacity and ChannelTransactionCapacity Properties

In this blog post we'll describe the different properties and behavior of a channel's capacity or batch size
This post is originally from Jeff Lord, posted on flume users (FLUME-1829)

Batch Size

- When configured by client code using the flume-core-sdk , to send events to flume avro source.
The flume client sdk has an appendBatch method. This will take a list of events and send them to the source as a batch. This is the size of the number of events to be passed to the source at one time.
- When set as a parameter on HDFS sink (or other sinks which support BatchSize parameter)
This is the number of events written to file before it is flushed to HDFS

Channel Capacity

- This is the maximum capacity number of events of the channel.

Channel Transaction Capacity

- This is the max number of events stored in the channel per transaction.

How will setting these parameters to different values, affect throughput, latency in event flow?

In general you will see better throughput by using memory channel as opposed to using file channel at the loss of durability.

- The channel capacity is going to need to be sized such that it is large enough to hold as many events as will be added to it by upstream agents. Ideal flow would see the sink draining events from the channel faster than it is having events added by its source.
- The channel transaction capacity will need to be smaller than the channel capacity.
 - If your Channel capacity is set to 10000 then Channel Transaction Capacity should be set to something like 100.
- Specifically if we have clients with varying frequency of event generation, i.e. some clients generating thousands of events/sec, while others at a much slower rate, what effect will different values of these params have on these clients ?
 - Transaction Capacity is going to be what throttles or limits how many events the source can put into the channel. This going to vary depending on how many tiers of agents/collectors you have setup. In general though this should probably be equal to whatever you have the batch size set to in your client.

With regards to the hdfs batch size, the larger your batch size the better performance will be. However, keep in mind that if a transaction fails the entire transaction will be replayed which could have the implication of duplicate events downstream.