

IBM MQ for z/OS v9.1.2 Active log acceleration with zHyperWrite

[tonySharkey](#)

Published on 26/03/2019

IBM MQ for z/OS version 9.1.2 introduces support for zHyperWrite on active log data sets.

Using zHyperWrite on MQ's active log datasets that are already replicated using IBM Metro Mirror can offer an improvement in the rate that the MQ log task is able to sustain.

Improved MQ active log performance is important for a number of reasons:

- Allows more work to get done with the same hardware footprint
- Better able to handle workload spikes
- Reduces costs

With persistent messaging workloads, the rate at which MQ is able to process the workload is largely dependent upon the rate at which the logger task is able to complete its I/O. In turn, this rate is almost entirely dependent on the rate at which the I/O subsystem is able to respond to requests.

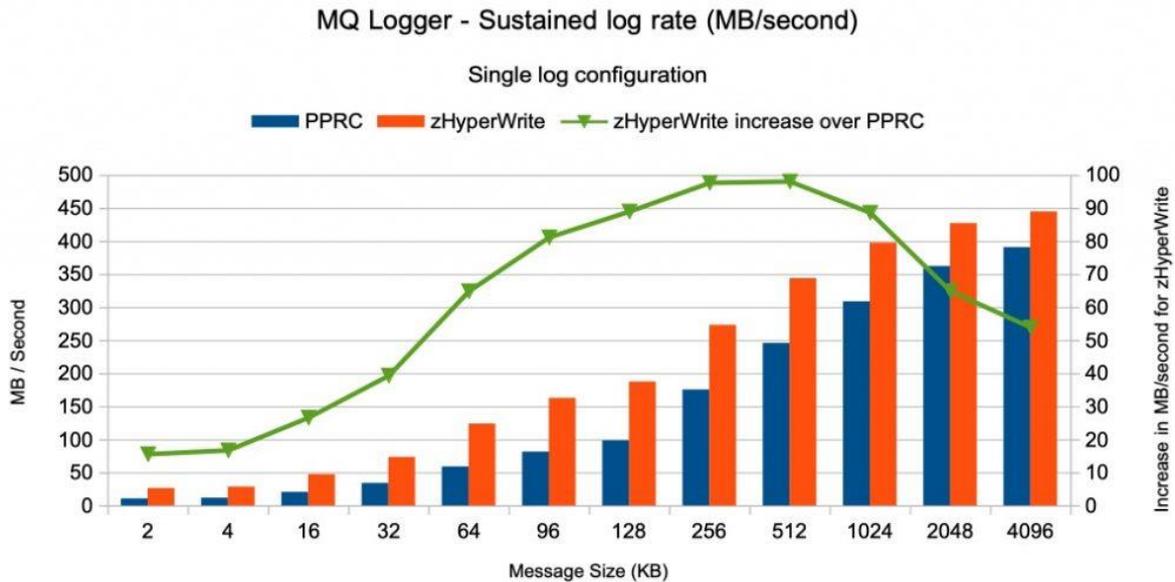
In this environment, one way to alleviate this constraint is to reduce the I/O times.

When running synchronous replication technologies such as Metro Mirror (PPRC), the I/O times can be significantly higher than when running without datasets replicated. The zHyperWrite technology can significantly reduce the I/O times when replicating datasets.

In our measurements, zHyperWrite benefited our workloads in a number of ways:

- Reduced I/O times by **up to 60%**.
- Reduced elapsed time for commit by **up to 60%**, which can reduce contention.
- Improved the sustained log rate, allowing each queue manager to process **up to 2.4 times** the volume of workload.

The chart below shows the improvement in the maximum sustained log rate for a range of message sizes from 2KB to 4MB on our performance test systems.



Maximum sustained rate of MQ Logger with single logs – comparing PPRC with zHyperWrite on IBM z14 and DS8870

The chart shows that with a workload using 256KB messages driving the logger task to capacity, we saw an increase in log rate of over 90MB / second.

The preferred configuration of dual logging saw the performance benefits of zHyperWrite reduce due to our system constraints. As we have only a single DS8870 for performance test, both the primary and secondary (mirrored) copies of the active logs were written to the same device. When driving high volumes of log data through the system, totalling 1.5GB/second, we saw our non-volatile storage being saturated, and as a result saw an increase in the proportion of I/Os impacted waiting for disk cache.

Where can I find more performance information?

For the in-depth performance report, see [“zHyperWrite MQ Active Log Acceleration”](#)

How do I enable zHyperWrite with MQ

Provided your active logs are replicated using Metro Mirror, previously known as Synchronous Peer to Peer Remote Copy (PPRC), to use zHyperWrite with the active logs, you need to:

- Configure IBM MQ to use zHyperWrite, and
- The active logs need to be on zHyperWrite capable volumes

If both of these conditions are met, writes to active logs are enabled for zHyperWrite.

You can configure IBM MQ to use zHyperWrite by using one of the following methods:

- Specify ZHYWRITE(YES) in the system parameter module.
- Issue the command SET LOG ZHYWRITE(YES).

Ensure the following conditions are true for active log data sets to be on zHyperWrite capable volumes:

- Enable the volumes for Metro Mirror, and the volumes support zHyperWrite
- Ensure that the volumes are HyperSwap enabled
- Specify HYPERWRITE=YES in the IECIOSxx parameter

If all the preceding conditions are met, then writes to the active logs are enabled for zHyperWrite.

If one, or more, of these conditions are not met, IBM MQ writes to the active logs as normal, and Metro Mirror replicates the writes if it is configured.

Finally

For this and other performance reports, keep an eye on the [MQ performance github page](#).