

Recommendations for Nodes and Encoding

While the cluster will fail a `WRITE` request in a cluster that does not have at least $k+p$ Netmail Store nodes available, the minimum number of required nodes is actually higher, to allow space for recovery in the event of a failed volume.

Number of nodes

- The *minimum* number of nodes to ensure base functionality is $k+(p*2)$.
- The *recommended* number for performance and load distribution is $(k+p)*2$.
- 5:2 encoding requires at least 9 Netmail Store nodes, preferably 14 nodes.

RAM per node

Nodes using erasure coding will benefit from additional CPU cores (and Netmail Store multi-server processes to utilize them).

The recommended minimum of RAM per node is 4 GB.

Encoding level

- Do *not* run Netmail Store using an EC level with *fewer than 2 parity segments*, except in (rare) instances with transient data that easily can be reproduced or regenerated.
- Good baseline encoding is **4:2** or **5:2**; this ensures data protection equal to or better than replication, with a smaller footprint.
- Subclusters have specific encoding considerations.

Performance considerations

Write throughput for large objects that are erasure-coded is approximately equivalent to replicated objects written with Replicate on Write. Smaller objects are less performant due to the overhead of creating the parity segments relative to the time to write only the content body.

- **Parity Count.** The biggest impact on performance for similar-sized objects is the parity count: the smaller the number of required parity segments, the faster the write performance.
- **Segment Size.** If you can predict the size of the objects written to the cluster, you can increase the default segment size to greater than 200 MB to decrease the number of EC sets it takes to encode the object, thereby increasing performance. For example, for a 2 GB object with 5:2 encoding, increasing the segment size to 400 MB will ensure the object fits in one EC set.
- **Write Threads.** Erasure-coded objects require that more Netmail Store nodes participate in each write request. As a result, the number of client threads each node can service at a time is less than with replicated objects. For planning purposes, assume that each Netmail Store node can service 4 erasure-coding client write threads at a time before the cluster begins to return `503 Service Unavailable` responses, indicating the cluster is too busy to service more requests. That threshold may differ depending on the size of the cluster and the encoding you use.