

Enabling Erasure Coding

Netmail Store ships with erasure coding disabled by default. If licensed in your cluster, you can select the encoding level for erasure-coded objects for each individual object or all objects in the cluster.

Note: To obtain a license file with erasure coding enabled, contact your sales representative.

Erasure Coding by Cluster

To erasure code all cluster objects, set the following options in the cluster or node configuration file to enable the cluster to erasure code objects based on their size (in bytes):

- `ec.minStreamSize`. Sets the minimum size, in bytes, for an object to be automatically erasure-coded.
- `ec.encoding`. Sets the number of data and parity segments to be used when erasure coding objects (for example, 5 : 2).

Erasure Coding by Object

To erasure code individual objects, use the `?encoding=k:p` query arguments with an SCSP `WRITE` or `UPDATE` method or specify an encoding tuple in the lifepoints for the object (`reps=5:2`). This process erasure codes the object at the specified encoding level.

Adding the optional `erasureCoded=yes` argument will force the object to be erasure-coded, regardless of the cluster setting or object size.

EC Disk Footprint

The amount of disk space (or *footprint*) used for erasure-coded objects depends on the ratio of data to parity segments in the specified encoding. Use the following formula to roughly calculate the disk space that you can expect to see used by an EC object with one set of erasure segments:

$k+p \text{ total segments} / k \text{ data segments} * \text{ObjectSize} = \text{total object EC disk footprint}$

A 1 GB object with 5:2 encoding has this footprint: $(5 + 2) / 5 * 1 \text{ GB} = 1.4 \text{ GB}$ (vs. 3 GB for replication)

A 3 GB object with 5:2 encoding has this footprint: $(5 + 2) / 5 * 3 \text{ GB} = 4.2 \text{ GB}$ (vs. 9 GB for replication)

Note: Some system metadata is written with each EC segment, adding ~16 bytes per segment.

Methods Affected by Erasure Coding

These are the SCSP methods that are affected by erasure-coded objects:

Method	Description
WRITE	A new object written to the storage cluster is erasure-coded if it meets the EC criteria.
UPDATE	As with <code>WRITE</code> , if a non-erasure-coded object is updated and the data causes the object to meet the EC criteria, the object can be erasure-coded.
APPEND	As with <code>UPDATE</code> , if data appended to a non-EC object causes the object to meet the EC criteria, the object can be erasure-coded. APPENDs with version 6.0 are optimized to write a new set of EC segments with the appended data and update the manifest for a previously erasure-coded object, instead of rewriting the whole object to include the appended data as with replicated objects. Data appended to an object that was not previously erasure-coded can cause the object to become erasure-coded if the additional appended data pushes the object size over the configured <code>minStreamSize</code> threshold. Including an <code>?encoding=k:p</code> query argument with an <code>APPEND</code> will result in a <code>400 Bad Request</code> response. On an <code>APPEND</code> for an existing EC object, the new segments will be encoded with the parameters of the first existing erasure set.

COPY	Instead of rewriting the entire object to update the metadata, COPY for erasure-coded objects updates the metadata on the manifest only without rewriting any content data.
INFO	Returns the Manifest: ec header for an erasure-coded object.
READ	The cluster automatically retrieves all segments so a READ does not behave differently for an erasure-coded object.
DELETE	Netmail Store deletes both the manifests and the segments for the object.