Erasure Coding

What is Erasure Coding?

SwiftStack provides several protection schemes for data redundancy, including both replicas and Erasure Coding (EC), which are defined as part of a Storage Policy in SwiftStack. Both methods are extremely durable ways of protecting data in a cluster and can be applied to different sets of data with different storage needs.

EC is a method of protection in which data is broken into segments and are encoded with redundant data segments and stored across a set of different locations or storage media. As an alternative to RAID, it is a widely used mechanism to protect data in storage systems, such as SwiftStack. The goal of erasure coding is to enable the reconstruction of corrupted or lost due to some failure by using information about the data that's stored elsewhere in the cluster.

HIGHLIGHTS

- Protect data with erasure coding or replicas in the same cluster
- Erasure coding is more space-efficient than using replicas
- Amount of redundant data is a configurable property

USES

- Reduce the amount of storage space required for certain data sets
- Good for storing large files in an active archive

How is it different from storing multiple replicas?

With replicas, data is protected by storing more than one whole copy of the data. If you store a copy in a few different places (e.g. different hard drives in different servers), then you're protected against hardware failure causing data loss. If you store three copies of the data, then you can lose two drives containing that data at a time and still recover. This is how replicas works in SwiftStack, which provides excellent data availability, reliability and performance. Protecting data with replicas is especially good for small files that are frequently accessed. On the converse, storing multiple whole copies of a file will require more disk drive hardware than EC. When storing three copies of the data, for instance, the overhead is 200%.

EC in SwiftStack can be used for one data set while replicas are used for another. This is set on a bucket-by-bucket basis. While replicas uses more storage capacity, there is reduced overhead required to return to a policy-compliant state in the event of a failure. With EC, less storage hardware is required, but reconstructing the data is more CPU-intensive.

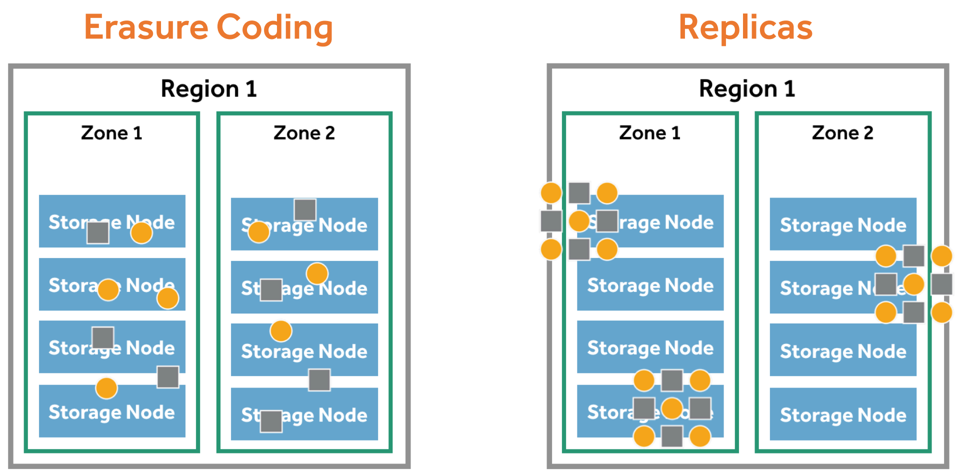

Erasure coding encodes and distributes parts of the data across the available drives, while **replicas** store whole copies of the data across the system.

Benefits of erasure coding

EC requires less disk drive hardware to store the same amount data as when protecting data with replicas. EC also have excellent performance characteristics when dealing with large objects, such as video. It is excellent at ingest and high-throughput streams due to the nature of breaking a large object into multiple segments and spreading those segments across many drives, which can participate in the parallel ingest process.

How is data protection implemented in erasure coding?

With EC, the data is broken into fragments and then a mathematical algorithm creates additional fragments, of which multiple combination can be used to recreate the data during a hardware failure. In mathematical terms, the protection offered by erasure coding can be represented in simple form by the following equation: $n = k + m$. The variable “k” is the original amount of data. The variable “m” stands for the extra or redundant data that are added to provide protection from failures. The variable “n” is the total number of data segments created after the erasure coding process. For instance, in an 8:4 EC configuration, four extra data segments (m) would be added to the 8 base segments (k). The 12 data fragments (n) would be spread across 12 drives, nodes or geographic locations. The original file could be reconstructed from 8 verified fragment and you only have a 50% overhead for drive space used.

To try SwiftStack for free, go to https://www.swiftstack.com/try-it-now/.

For additional assistance or to learn more, always feel free to contact us. We're here to help.

Phone - (415) 625-0293
Email - contact@swiftstack.com
Chat - Just go to swiftstack.com and look for the chat pop-up in the bottom right