FORM 2

THE PATENTS ACT, 1970
(39 of 1970)
AND
THE PATENTS RULES, 2003

COMPLETE SPECIFICATION

(See Section 10; rule 13)

TITLE OF THE INVENTION

"METHOD AND APPARATUS FOR ALLOCATING ERASURE CODED DATA TO DISK STORAGE"

APPLICANT

SimpliVity Corporation
of 8 Technology Drive, Westborough, Massachusetts 01581-1756, USA; Nationality: USA

The following specification particularly describes
the invention and the manner in which
it is to be performed

Clean Copy of Claims

1. A method comprising:

allocating encoded data objects for storage in a pool of disk storage devices comprising multiple disk drives;

at least some of the encoded objects being generated by different error correction algorithms and allocation is not restricted to a group of drives based on the algorithm utilized to generate the encoded object;

each disk drive is divided into multiple logical storage units, each defined as a partition;

a logical storage unit group (LSU group) comprising a group of partitions extending across a group of disk drives, wherein the partitions of the LSU group are each of the same size and cannot be on the same drive;

the method including generating an allocation bitmask to describe how the encoded data object is to be stored;

the allocating of each encoded data object being performed by comparing the allocation bitmask for the encoded data object and a single allocation bitmap that extends across the LSU group to identify available locations for storage of the object, wherein the object is allocated to one available location defined by contiguous bits of the bitmap, or to multiple locations aligned along a common partition boundary allowing a gap of non-contiguous allocation between the two locations.

2. The method of claim 1, wherein:

the allocating step includes allocating the encoded objects on different drives.

3. The method of claim 2, wherein:

the allocating step includes allocating encoded objects encoded by different error correction algorithms.

4. The method of claim 1, wherein:

the method includes encoding data by different error correction algorithms to generate the encoded objects and maintaining an index which maps the encoded object to its respective error correction algorithm.

5. The method of claim 1, wherein:
the allocating step includes allocating multiple encoded objects on the same logical storage unit.

6. The method of claim 1, wherein:
the allocating step includes allocating multiple encoded objects on the same logical storage unit group.

7. The method of claim 1, including:
using the allocation bitmap to request allocation units aligned with a common partition boundary.

8. The method of claim 1, wherein:
the object size of the data being encoded is fixed.

9. The method of claim 1, wherein:
the object size of the data being encoded is variable.

10. The method of claim 1, wherein:
the data objects are encoded by different categories of error correction algorithms.

11. The method of claim 1, including:
providing an index of the encoded data objects which maps each encoded data object to its respective error correction algorithm.

12. The method of claim 1, wherein:

the allocating step includes using an allocation bitmap assigning multiple allocation bits per chunk group of the respective error correction algorithm.

13. The method of claim 12, wherein:
the allocation bitmap maps to a logical address space.

14. The method of claim 13, wherein:
a logical object number (LON) defines a pointer to the encoded object.

15. The method of claim 1, wherein:
a pointer to the encoded object is stored in an index record.

16. The method of claim 15, wherein:
the index record includes multiple pointers to the encoded object.

17. The method of claim 12, wherein the allocating step uses a boundary bitmap marking the allocation unit for an initial chunk of the encoded object.

18. A computer-readable medium having stored thereon instructions which perform, when loaded into a computer, the method steps according to claim 1.

19. A programmable logic configured to implement the method steps according to claim 1.

20. A data storage system comprising:
an error correction algorithm component operable to select, for different incoming data objects, different error correction algorithms for generating encoded data objects; and
a disk storage allocation component for allocating the encoded data objects for storage in a pool of disk storage devices comprising multiple disk drives, at least some of the encoded objects generated by different error correction

algorithms being allocated to the same or a different group of drives that is not based on the algorithm utilized to generate the encoded object;

wherein each disk drive is divided into multiple logical storage units, each defined as a partition;

a logical storage unit group (LSU group) comprising a group of partitions extending across a group of disk drives wherein the partitions of the LSU group are of each of the same size and cannot be on the same drive;

the allocation component comparing an allocation bitmask describing how the encoded data object is to be stored to a single allocation bitmap that extends across the LSU group to identify available locations for storage of the object, wherein the object is allocated to one available location defined by contiguous bits of the bitmap, or to multiple locations aligned along a common partition boundary allowing a gap of non-contiguous allocation between the two locations.

21. The storage system of claim 20, further comprising:
the pool of disk storage devices for storing the encoded data objects.

22. The storage system of claim 20, wherein:
the allocation bitmap assigns multiple bits per chunk group of the respective error correction algorithm.

23. The storage system of claim 20, including:
an index of the encoded data objects which maps each encoded data object to its respective error correction algorithm.

24. In a computing environment for locating data storage, a data structure comprising an allocation bitmap to request available allocation units for storing encoded data objects in a pool of disk storage devices comprising multiple disk drives;

at least some of the encoded objects being generated by different error correction algorithms and allocation is not restricted to a group of drives based on the algorithm utilized to generate the encoded object;

each disk drive is divided into multiple logical storage units, each defined as a partition;

a logical storage unit group (LSU group) comprising a group of partitions extending across a group of disk drives, wherein partitions of the LSU group are each of the same size and cannot be on the same drive;

the single allocation bitmap extends across the LSU group to identify, by a comparison of an allocation bitmask describing how an encoded data object is to be stored, to the allocation bitmap, available locations for storage of the object, the bitmap including multiple available locations aligned along a common boundary allowing a gap of non-contiguous allocation between the two locations.

Dated this 16 day of June 2014

Arindam Paul
REG.NO:IN/PA-174
of Depenning & Depenning
Agent for the Applicants