the resiliency set may be a partial data redundancy. In other words, data of the write operation may be spread out across N memory regions in such a way that the data of the write operation may be recovered from at least M memory regions, where M is less than N. In one such example, the relationship between the memory regions of the resiliency set is an erasure-coding-based logical relationship.

The computer readable storage medium my include instructions executable to form a resiliency logic cluster that identifies a logical cluster of resiliency logics corresponding to the memory regions of the resiliency set.

A client device may be provided that includes a processor that is configured to: write data to a first region of memory on a first memory appliance in response to a write operation, wherein the first region of memory is external memory to the client device, and wherein the first region of memory is accessible by the client device over a network via client-side memory access in which a first communication interface of the first memory appliance is configured to access the first region of memory on the first memory appliance; and cause the data of the write operation to be written to a second region of memory on a second memory appliance, wherein a resiliency set comprises a plurality of memory regions including the first region of memory and the second region of memory, and wherein the data of the write operation is recoverable from a subset of the memory regions, the subset excluding the first region of memory or the second region of memory.

For example, the processor may be configured to write the data of the write operation to the second region of memory on the second memory appliance via client-side memory access. Alternatively, the processor may be configured to transmit a notification message over the network to the second memory appliance and/or a second client device, in response to which, the data of the write operation is written to the second region of memory on the second memory appliance.

The processor may be configured to form a snapshot of memory by removal of a memory region from the resiliency set, wherein the snapshot includes contents of the memory region removed from the resiliency set.

In some examples, the processor is configured to notify an observer logic for the second region of memory that a multi-write is upcoming, and to notify the observer logic if the multi-write completes.

The processor may be configured to rebuild resiliency of the memory regions in the resiliency set asynchronously of the write operation.

A second action may be said to be "in response to" a first action independent of whether the second action results directly or indirectly from the first action. The second action may occur at a substantially later time than the first action and still be in response to the first action. Similarly, the second action may be said to be in response to the first action even if intervening actions take place between the first action and the second action, and even if one or more of the intervening actions directly cause the second action to be performed. For example, a second action may be in response to a first action if the first action sets a flag and a third action later initiates the second action whenever the flag is set.

To clarify the use of and to hereby provide notice to the public, the phrases "at least one of <A>, <B>, . . . and <N>" or "at least one of <A>, <B>, . . . <N>, or combinations thereof" or "<A>, <B>, . . . and/or <N>" are defined by the Applicant in the broadest sense, superseding any other implied definitions hereinbefore or hereinafter unless expressly asserted by the Applicant to the contrary, to mean one or more elements selected from the group comprising A, B, . . . and N. In other words, the phrases mean any combination of one or more of the elements A, B, . . . or N including any one element alone or the one element in combination with one or more of the other elements which may also include, in combination, additional elements not listed.

While various embodiments of the innovation have been described, it will be apparent to those of ordinary skill in the art that many more embodiments and implementations are possible within the scope of the innovation. Accordingly, the innovation is not to be restricted except in light of the attached claims and their equivalents.

What is claimed is:

1. A method comprising:

writing, by a client device over a network via client-side memory access, a first data to a first region of memory of a first memory appliance in response to a write operation, wherein the first data is derived from data of the write operation, wherein the first region of memory is external primary memory, which is primary memory to the client device but external to the client device, wherein the client comprises a processor, wherein primary memory is any memory in an address space addressable by the processor, and wherein a first communication interface of the first memory appliance is configured to access the first region of memory on the first memory appliance for any client-side memory access of the first region of memory; and

writing, by the client device via client-side memory access, a second data to a second region of memory of a second memory appliance, wherein the second data is derived from the data of the write operation, wherein the second region of memory is external primary memory to the client device, wherein a resiliency set comprises a plurality of memory regions including the first region of memory and the second region of memory, and wherein the data of the write operation is recoverable from a subset of the memory regions, the subset excluding the first region of memory or the second region of memory.

2. The method of claim **1**, wherein said writing to the first region of memory completes before said writing to the second region of memory starts.

3. The method of claim **1**, wherein said writing to the first region of memory and said writing to the second region of memory are performed in parallel.

4. The method of claim **1** further comprising waiting until the first data and the second data are written to the first region of memory and the second region of memory, respectively, before completing the write operation.

5. The method of claim **1**, wherein said writing to the first region of memory triggers a page fault with the communication interface of the first memory appliance and a corresponding page fault hander completes handling of the page fault after waiting for one or more conflicting resiliency rebuild operations to finish.

6. The method of claim **1** further comprising transmitting a notification message from the client device to a second client device, wherein a third data is written to a third region of memory of a third memory appliance in response to the notification message, the third data derived from the data of the write operation, the third region of memory included in the memory regions of the resiliency set.

7. A non-transitory computer readable storage medium comprising computer executable instructions, the computer

executable instructions executable by a processor, the computer executable instructions comprising:

instructions executable to write a first data from a first client device to a first region of memory of a first memory appliance via client-side memory access in response to a write operation received at the first client device, wherein the first data is derived from data of the write operation, wherein the first region of memory is external primary memory, which is allocatable as primary memory of the first client device but is external to the first client device, wherein primary memory is any memory in an address space addressable by the processor, and wherein a first communication interface of the first memory appliance is configured to access the first region of memory on the first memory appliance for any client-side memory access of the first region of memory; and

instructions executable to transmit a notification message to a second client device and/or a second memory appliance, the notification message indicating a second data is to be written to a second region of memory of the second memory appliance, wherein the second data is derived from the data of the write operation, wherein the second region of memory is external primary memory, which is allocatable as primary memory of the first client device or the second client device but is external to the first client device and the second client device, wherein a resiliency set comprises a plurality of memory regions including the first region of memory and the second region of memory, and wherein the data of the write operation is recoverable from a subset of the memory regions, the subset excluding the first region of memory or the second region of memory.

8. The non-transitory computer readable storage medium of claim **7**, wherein the notification message identifies the second data.

9. The non-transitory computer readable storage medium of claim **7** further comprising instructions executable to form a snapshot of primary memory by removal of a memory region from the resiliency set, wherein the snapshot includes contents of the memory region removed from the resiliency set.

10. The non-transitory computer readable storage medium of claim **9** further comprising instructions executable to wait for conflicting operations to complete upon notification of a start of the snapshot.

11. The non-transitory computer readable storage medium of claim **7**, wherein a relationship between the memory regions of the resiliency set is that each of the memory regions includes a corresponding copy of data common to all of the memory regions, and wherein the first data is the same as the second data, and the data of the write operation is included in the first data and the second data.

12. The non-transitory computer readable storage medium of claim **7**, wherein a relationship between the memory regions of the resiliency set is a partial data redundancy.

13. The non-transitory computer readable storage medium of claim **7**, wherein a relationship between the memory regions of the resiliency set is an erasure-coding-based logical relationship.

14. The non-transitory computer readable storage medium of claim **7** further comprising instructions executable to form a resiliency logic cluster that identifies a logical cluster of resiliency logics corresponding to the memory regions of the resiliency set.

15. A client device comprising:

a processor configured to:

write a first data to a first region of memory on a first memory appliance in response to a write operation, wherein the first data is derived from data of the write operation, wherein the first region of memory is external primary memory to the client device, wherein primary memory is any memory in an address space addressable by the processor and wherein the first region of memory is accessible by the client device over a network via client-side memory access in which a first communication interface of the first memory appliance is configured to access the first region of memory on the first memory appliance; and

cause a second data to be written to a second region of memory on a second memory appliance, wherein the second data is derived from the data of the write operation, wherein a resiliency set comprises a plurality of memory regions including the first region of memory and the second region of memory, and wherein the data of the write operation is recoverable from a subset of the memory regions, the subset excluding the first region of memory or the second region of memory.

16. The client device of claim **15**, wherein the processor is configured to write the second data to the second region of memory on the second memory appliance via client-side memory access.

17. The client device of claim **15**, wherein the processor is configured to transmit a notification message over the network to the first memory appliance, the second memory appliance, and/or a second client device, in response to which, the second data is written to the second region of memory on the second memory appliance.

18. The client device of claim **17**, wherein the processor is configured to form a snapshot of memory by removal of a memory region from the resiliency set, wherein the snapshot includes contents of the memory region removed from the resiliency set.

19. The client device of claim **15**, wherein the processor is configured to notify an observer logic for the second region of memory that a multi-write is upcoming, and to notify the observer logic if the multi-write completes.

20. The client device of claim **15**, wherein the processor is configured to rebuild resiliency of the memory regions in the resiliency set asynchronously of the write operation.

* * * * *