ered. In one embodiment, the instructions are stored on a removable media device for reading by local or remote systems. In other embodiments, the logic or instructions are stored in a remote location for transfer through a computer network or over telephone lines. In yet other embodiments, the logic or instructions are stored within a given computer, CPU, GPU, or system.

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 non-transitory computer-readable storage medium comprising a plurality of instructions executable with a processor, the instructions including:

instructions executable to receive, at a client device from a component in the client device, a request to allocate a portion of memory, wherein a memory appliance and the client device are connected by a network, and a region of memory in the memory appliance is allocated as external memory prior to receipt of the request to allocate the portion of memory at the client device, wherein external memory is memory that is external to the client device but is primary memory to the client device;

instructions executable to select, at the client device and in response to the request to allocate the portion of memory, a subset of the region of memory in the memory appliance; and

instructions executable to map, at the client device, at least the portion of memory to a virtual address space, wherein data in the portion of memory is accessible over a network via client-side memory access in which a communication interface of the memory appliance is configured to access the region of memory in the memory appliance.

2. The computer-readable storage medium of claim **1**, wherein the instructions executable to map the at least the portion of memory to the virtual address space are executable to map all of the region of memory to the virtual address space.

3. The computer-readable storage medium of claim **2**, wherein the instructions executable to map the at least the portion of memory to the virtual address space are executable to map the at least the portion of memory after the subset of the region of memory is selected.

4. The computer-readable storage medium of claim **2**, wherein the instructions executable to map the at least the portion of memory to the virtual address space are executable to map the at least the portion of memory before the subset of the region of memory is selected.

5. The computer-readable storage medium of claim **1**, wherein the instructions executable to select the subset of the region of memory are invoked through an application-level memory allocator.

6. The computer-readable storage medium of claim **1**, wherein the region of memory in the memory appliance is accessible by the client device over the network via client-side memory access from an operating system, a kernel, a device driver, a virtual machine, a hypervisor, a container, and/or a jail included in the client device.

7. The computer-readable storage medium of claim **1**, wherein the external memory is accessible from via a memory swapping interface, a swap device, and/or a swap file included in the client device.

8. A method comprising:

receiving a memory allocation request at a client device from a component of the client device;

selecting, at the client device and in response to the memory allocation request, a subset of a region of memory in a memory appliance to be a portion of memory allocated at the client device, wherein the client device and the memory appliance are in communication over a network, and wherein the region of memory in the memory appliance is memory allocated for the client device before the memory allocation request is received; and

mapping, at the client device, at least the portion of memory to a virtual address space, wherein data in the portion of memory is accessible via a client-side memory access, wherein a communication interface of the memory appliance is configured to access the subset of the region of memory in the memory appliance, as part of the client-side memory access, over an interconnect logically between the communication interface and a memory controller of the memory appliance.

9. The method of claim **8**, wherein the selecting is performed by logic operating within a virtualization instance.

10. The method of claim **8**, wherein the client device is a first client device, the method further comprising accessing the data in the portion of memory over the network via client-side memory access by a second client device.

11. The method of claim **8** further comprising causing the data received from the memory appliance to be decrypted and/or causing the data to be encrypted prior to transmission to the memory appliance.

12. The method of claim **8**, wherein the client device is a first client device, the first client device includes a virtual machine, the subset of the region of memory is allocated to the virtual machine, the method further comprising migrating the virtual machine from the first client device to a second client device without copying the data by accessing the data in the subset of the region of memory by the virtual machine from the second client device.

13. The method of claim 8, wherein the client device is a first client device, the first client device includes a virtual machine, the subset of the region of memory is allocated to the virtual machine, the method further comprising migrating the virtual machine from the first client device to a second client device without copying the data by discarding a cached copy of the data from the virtual machine and subsequently accessing the data in the subset of the region of memory by the virtual machine from the second client device.

14. A system comprising:

a client device comprising a first processor; and

a memory appliance comprising:

a communication interface configured to communicate over a network,

a memory comprising a region of memory allocated as external memory for the client device, and

a second processor, wherein the external memory is memory that is external to the client device but primary memory to the client device, wherein the first processor of the client device is configured to allocate a subset of the region of memory as a portion of external memory in response to a memory allocation request, and wherein data in the portion of external memory is accessible by the client device via client-side memory access through the communication interface, the communication interface configured to access the subset of the region of memory in response to the client-side memory access of the data in the portion of external memory.

15. The system of claim 14, wherein the first processor is configured to encrypt the data before the data is written to a backing store, and/or the first processor is configured to decrypt the data read from the backing store.

16. The system of claim 14, wherein the first processor is further configured to map at least the portion of the external memory to a virtual address space of the client device.

17. The system of claim 14, wherein the client device includes a virtual machine and the external memory is for use by the virtual machine.

18. The system of claim 14, wherein the memory allocation request is a second memory allocation request and the memory comprises a region allocation logic executable by the second processor to allocate the region of memory as external memory for the client device in response to a first memory allocation request, and wherein the region allocation logic executes within a virtualization instance.

19. The system of claim 14, wherein the memory allocation request is a second memory allocation request and the memory comprises a region allocation logic executable by the second processor to allocate the region of memory as external memory for the client device in response to a first memory allocation request, and wherein the region allocation logic executes within a hypervisor.

20. The system of claim 14, wherein the memory allocation request is a second memory allocation request and the memory comprises a region allocation logic executable by the second processor to allocate the region of memory as external memory for the client device in response to a first memory allocation request, and wherein the region allocation logic executes within a cloud computing interface.

21. The system of claim 14, wherein a second client executes within a virtualization instance of the memory appliance and wherein the second client allocates memory of a second memory appliance.

22. The system of claim 14, wherein a second client operates within a hypervisor of the memory appliance and wherein the second client allocates memory of a second memory appliance.

23. The system of claim 14, wherein the communication interface, and/or the second processor is included in a FPGA (Field-Programmable Gate Array) and/or an ASIC (Application-specific Integrated Circuit).

* * * * *