puter-readable media, which includes CD-ROMs, volatile or non-volatile memory such as ROM and RAM, or any other suitable storage device. As another example, the systems, components, and flow diagrams illustrated herein include one or more components and/or operations. Each system, component, or set of operations may include fewer, additional, or different components or operations. Components or operations indicated as optional in the drawings may just be one example of components or operations that are optional. Components or operations not marked as optional may, in fact, be optional in some examples.

Furthermore, although specific components of innovations were described, methods, systems, and articles of manufacture consistent with the innovation may include additional or different components. For example, a processor may be implemented as a microprocessor, microcontroller, application specific integrated circuit (ASIC), discrete logic, or a combination of other type of circuits or logic. Similarly, memories may be DRAM, SRAM, Flash or any other type of memory. Flags, data, databases, tables, entities, and other data structures may be separately stored and managed, may be incorporated into a single memory or database, may be distributed, or may be logically and physically organized in many different ways. The components may operate independently or be part of a same program. The components may be resident on separate hardware, such as separate removable circuit boards, or share common hardware, such as a same memory and processor for implementing instructions from the memory. Programs may be parts of a single program, separate programs, or distributed across several memories and processors.

The respective logic, software or instructions for implementing the processes, methods and/or techniques discussed throughout this disclosure may be provided on computer-readable media or memories or other tangible media, such as a cache, buffer, RAM, removable media, hard drive, other computer readable storage media, or any other tangible media or any combination thereof. The tangible media include various types of volatile and nonvolatile storage media. The functions, acts or tasks illustrated in the figures or described herein may be executed in response to one or more sets of logic or instructions stored in or on computer readable media. The functions, acts or tasks are independent of the particular type of instructions set, storage media, processor or processing strategy and may be performed by software, hardware, integrated circuits, firmware, micro code, or any type of other processor, operating alone or in combination. Likewise, processing strategies may include multiprocessing, multitasking, parallel processing and/or any other processing strategy known now or later discovered. 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 first client device from a component in the first client device, a request to allocate a portion of memory, wherein a memory appliance and the first client device are on a network, and a region of memory of the memory appliance is allocated as external memory prior to receipt of the request to allocate the portion of memory at the first client device, wherein the external memory is external to the first client device and a second client device, but is primary memory to the first client device and the second client device;

instructions executable to select, by a first client logic within the first client device, in coordination with a second client logic within the second client device and in response to the request to allocate the portion of memory, a subset of the region of memory of the memory appliance, the first client logic being otherwise distinct from the second client logic aside from coordinating with the second client logic in using the memory of the memory appliance; and

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

wherein selection of the subset of the region of memory in coordination with the second client device is based on a memory-allocation data structure, wherein the memory-allocation data structure is shared with the second client device.

2. The computer readable medium of claim **1** further comprising instructions executable to update the memory-allocation data structure to indicate that the portion has been allocated as part of the selection of the subset of the region of memory in coordination with the second client device.

3. The computer readable medium of claim **2**, wherein the memory-allocation data structure is included in metadata associated with the region of memory.

4. The computer readable medium of claim **1**, wherein the memory-allocation data structure is accessible from the first client device and the second client device via client-side memory access.

5. The computer readable medium of claim **1** comprising instructions executable to access the memory-allocation data structure atomically as part of the selection of the subset of the region of memory in coordination with the second client device.

6. The computer readable medium of claim **1**, wherein access to the memory-allocation data structure is coordinated based on one or more synchronization primitives.

7. The computer readable medium of claim **1**, wherein the memory-allocation data structure is included in the region of memory of the memory appliance.

8. The computer readable medium of claim **1**, wherein the memory-allocation data structure includes one or more indicators that indicate whether corresponding portions of the region are allocated.

9. The computer readable medium of claim **1**, wherein selection of the subset of the region of memory in coordination with the second client device is based on message-passing between the first client and the second client.

10. The computer readable medium of claim **1**, wherein the at least the subset of the region of memory mapped to the virtual address space includes all of the region of memory and wherein the at least the subset of the region of memory is mapped to the virtual address space before the request to allocate a portion of memory is received and/or before the subset of the region of memory is selected.

11. The computer readable medium of claim **1**, wherein a second subset of the region of memory is allocable by the second client device in response to a request to allocate a second portion of memory.

12. The computer readable medium of claim **1** further comprising instructions executable to initialize at least a second portion of external memory after the at least the second portion of external memory is mapped at the first client device.

13. The computer readable medium of claim **1** further comprising instructions executable to free the subset of the region of memory in response to destruction of a corresponding application logic or to a determination that the corresponding application logic is destroyed.

14. A method comprising:

receiving, at a first client device, a request to allocate a portion of memory, wherein a memory appliance and the first client device are connected by a network, and a region of memory of the memory appliance is allocated as external memory before receipt of the request to allocate the portion of memory at the first client device;

selecting, by a first client logic within the first client device and in coordination with a second client logic within a second client device on the network, a subset of the region of memory for the portion of memory allocated, wherein the portion of memory allocated is external to the first client device but is primary memory to the first client device, wherein by coordination with the second client device, the subset of the region of memory is selected without the second client device having selected the same subset of the region of memory for allocation as external memory for the second client device; and wherein the first client logic is otherwise distinct from the second client logic aside from coordinating with the second client logic in using the memory of the memory appliance; and

mapping, at the first client device, at least the subset of the region of memory to a virtual address space,

wherein selecting the subset of the region of memory in coordination with the second client device comprises accessing a memory-allocation data structure that is shared by the first client device and the second client device, the memory-allocation data structure indicating which portions of the region of memory are allocated by a memory user of the first client device.

15. The method of claim **14**, wherein selecting the subset of the region of memory in coordination with the second client comprises identifying a free portion of the region of memory of the memory appliance from the memory-allocation data structure shared by the first client and the second client.

16. The method of claim **14** further comprising migrating an application logic from the first client device to the second client device by mapping, at the second client, the portion of the region of memory of the memory appliance that is indicated by the memory-allocation data structure to be allocated to the application logic on the first client.

17. The method of claim **16**, wherein migrating the application logic includes migrating the application logic from the first client device to the second client device without copying any data of the allocated portion of memory from the first client device to the second client device.

18. The method of claim **16** further comprising determining the portion of the region of memory of the memory appliance allocated to the application logic on the first client by accessing the memory-allocation data structure from the second client device.

* * * * *