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 computer executable instructions, which when executed by a processor, perform the following steps:

receiving a first request for an external primary memory allocation, the first request received from an application logic unit of a device, wherein the first request for the external primary memory allocation is a request for memory that is external to the device but that is primary memory to the device;

determining an allocation strategy for the external primary memory allocation in response to receipt of the first request, the allocation strategy comprising identification of a memory appliance on which to allocate a region of memory;

allocating the region of memory by sending, over a network via a communication interface, a second request for allocation of the region of memory on the identified memory appliance;

receiving, at the device, a request to allocate a slab of the external primary memory, the request to allocate the slab received from the application logic unit of the device;

selecting, independent of the memory appliance, at the device in response to the request to allocate the slab, a subset of the previously allocated region of memory to be the slab; and

mapping, at the device, the slab of the external primary memory into a virtual address space, wherein data in the slab is accessible over the communication interface with client-side memory access independent of a central processing unit of the identified memory appliance.

**2**. The non-transitory computer-readable storage medium of claim **1**, wherein determining the allocation strategy includes selecting the memory appliance from a plurality of memory appliances on which to allocate the region of memory.

**3**. The non-transitory computer-readable storage medium of claim **1**, wherein determining the allocation strategy comprises determining a number of regions to allocate for the external primary memory allocation.

**4**. The non-transitory computer-readable storage medium of claim **3**, wherein the memory appliance is a first memory appliance, and determining the allocation strategy comprises:

selecting the first memory appliance and a second memory appliance from a plurality of memory appliances, and

allocating a first region on the first memory appliance and a second region on the second memory appliance.

**5**. The non-transitory computer-readable storage medium of claim **1**, wherein determining the allocation strategy comprises provisioning a predetermined amount of the external primary memory for the device from which memory is allocated to the device in fulfillment of allocation requests received after receipt of the first request for the external primary memory allocation, wherein the allocation requests include the request to allocate the slab of the external primary memory.

**6**. The non-transitory computer-readable storage medium of claim **1**, wherein determining the allocation strategy comprises determining a distribution across a plurality of memory appliances of memory to allocate for the external primary memory allocation based on a network location of the device relative to the memory appliances, the memory appliances comprising the memory appliance.

**7**. The non-transitory computer-readable storage medium of claim **1**, wherein determining the allocation strategy comprises determining a distribution across memory appliances of memory to allocate for the external primary memory allocation based on a network bandwidth between the device and the memory appliances, the memory appliances comprising the memory appliance.

**8**. The non-transitory computer-readable storage medium of claim **1**, wherein determining the allocation strategy comprises determining an amount of memory to allocate for the external primary memory allocation based on a profile associated with the device.

**9**. The non-transitory computer-readable storage medium of claim **1** further comprising computer executable instructions, which when executed with the processor, perform a step of receiving a request to create, destroy, or modify the external primary memory allocation.

**10**. An apparatus comprising:

an interconnect; and

a local memory comprising:

an application logic unit; and

a client logic unit configured to cause, in response to a memory allocation request, allocation of a region of an external memory that is accessible on a memory appliance by the apparatus over the interconnect, wherein the external memory is memory that is external to the apparatus but primary memory to the apparatus, and

wherein the client logic unit is further configured to receive a request from the application logic unit to allocate a slab of the external memory,

wherein the client logic is further configured to select at the apparatus, independent of the memory appliance, a subset of the previously allocated region of the external memory to be the slab,

wherein the client logic is further configured to map, at the apparatus, the slab of the external memory into a virtual address space, and

wherein the client logic unit is further configured to access the external memory on the memory appliance over the interconnect with client-side memory access independent of a central processing unit of the memory appliance.

**11**. The apparatus of claim **10**, wherein the client logic unit is further configured to cache, in the local memory, portions of data stored in the slab of the external memory,

the data accessed in memory access operations performed on the slab of the external memory.

**12**. The apparatus of claim **10**, wherein the client-side memory access includes execution of a memory access operation that conforms to the Remote Direct Memory Access (RDMA) protocol.

**13**. The apparatus of claim **10**, wherein the local memory further comprises an allocation logic unit configured to receive the request from the client logic unit to allocate the external memory, wherein the allocation logic unit is further configured to select the memory appliance on which to allocate the region of the external memory, and the client logic unit is configured to access the region of the external memory over the interconnect with client-side memory access.

**14**. The apparatus of claim **13**, wherein the memory appliance is selected for allocation of the external memory from a plurality of memory appliances, the selection based on performance criteria of the memory appliances.

**15**. The apparatus of claim **10**, wherein the external memory is accessible by the application logic unit through a data interface included in the apparatus.

**16**. The apparatus of claim **15**, wherein the data interface is a memory allocation interface.

**17**. The apparatus of claim **15**, wherein the data interface is a memory swapping interface.

**18**. The apparatus of claim **15**, wherein the data interface is a memory caching interface.

**19**. The apparatus of claim **15**, wherein the data interface is a block-level interface.

**20**. The apparatus of claim **15**, wherein the data interface is a memory-mapped interface.

**21**. The apparatus of claim **15**, wherein the data interface is a graphics processor unit (GPU) accessible interface.

**22**. The apparatus of claim **15**, wherein the data interface is a character-level interface.

**23**. The apparatus of claim **15**, wherein the data interface is a hardware accessible interface.

**24**. The apparatus of claim **15**, wherein a peripheral device of the apparatus is provided access to the external memory via the data interface.

**25**. A memory appliance comprising:

a central processing unit;

a communication interface;

a memory; and

a region access unit configured to receive, over a network via the communication interface, a request to allocate a portion of the memory for use as an external primary memory of a client on the network, the external primary memory of the client is primary memory of the client that is external to the client, and the region access unit further configured to allocate the portion of the memory for the client;

wherein the communication interface is configured to provide the client access to the allocated portion of the memory on the memory appliance via client-side memory access in which the central processing unit of the memory appliance is bypassed; and

wherein the client is configured to select, independent of the memory appliance and in response to a request to allocate a slab of the external primary memory, a subset of the allocated portion of the memory on the memory appliance to be the slab, and to map the slab of the external primary memory into a virtual address space.

* * * * *