

117

While various embodiments have been described, it will be apparent to those of ordinary skill in the art that many more embodiments and implementations are possible. Accordingly, the embodiments described herein are examples, not the only possible embodiments and imple- 5
mentations.

What is claimed is:

1. A method comprising:
 - allocating, for an application logic, a region of external primary memory included in a memory appliance;
 - selecting, by a client device in response to a first request, a first portion of local primary memory in the client device to reclaim based on statistical data maintained for portions of the local primary memory;
 - determining a portion of external primary memory from the region of external primary memory;
 - copying data from the first portion of local primary memory to the portion of external primary memory; and
 - invalidating a portion of a first virtual address space at the client device corresponding to the first portion of the local primary memory.
2. The method of claim 1, wherein the first portion of local primary memory is selected by:
 - collecting, periodically or in response to an event or a condition, values stored in one or more portion tracking data structures into a collection, wherein each of the portion tracking data structures corresponds to a portion of local primary memory, and wherein values stored in the each of the portion tracking data structures represents the statistical data for each portion of local primary memory corresponding to the each of the portion tracking data structures; and
 - selecting the first portion of local primary memory based on the collection.
3. The method of claim 2, wherein:
 - the event comprises at least one of:
 - a page fault associated with the first virtual address space; or
 - a request to access the first virtual address space; and
 - the condition comprises at least one of:
 - running out of entries in the collection; or
 - running out of useful entries in the collection.
4. The method of claim 2, wherein:
 - the collection is a sorted collection which is sorted according to a reclaim strategy; or
 - after collecting the values stored in the one or more portion tracking data structures into the collection, the method further comprises:
 - sorting the values in the collection according to the reclaim strategy.
5. The method of claim 4, wherein the reclaim strategy comprises at least one of:
 - a least-recently-used strategy based on a timestamp or a generation counter; or
 - a not-frequently-used strategy based on a count of accesses or a count of page faults.
6. The method of claim 2, wherein selecting the first portion of local primary memory based on the collection comprises at least one of:
 - selecting a portion of local primary memory identified by an identifier associated with a minimum or maximum value in the collection as the first portion of local primary memory;
 - selecting a portion of local primary memory identified by an identifier associated with a value in the collection

118

- that is below or above a threshold value as the first portion of local primary memory; or
- randomly selecting one or more candidate portions of local primary memory, each of the one or more candidate portions mapping to a value in the collection, and selecting a portion of local primary memory from the one or more candidate portions of local primary memory as the first portion of local primary memory if the value corresponding to the portion of local primary memory is below or above the threshold value.
7. The method of claim 6, wherein the threshold value comprises at least one of:
 - a mean value of values stored in the collection;
 - a median value of the values stored in the collection;
 - a percentile value of the values stored in the collection;
 - a value calculated based on one or more of the values stored in the collection;
 - a value based on a timestamp; or
 - a value based on a generation counter.
8. The method of claim 2, wherein collecting the values stored in the one or more portion tracking data structures into the collection comprises:
 - collecting the values stored in the one or more portion tracking data structures by speculatively referencing entries in the one or more portion tracking data structures, wherein each of the entries corresponds to a portion of local primary memory and is able to be individually locked when referenced.
9. The method of claim 8, wherein in response to an entry in the one or more portion tracking data structures being already locked, skipping collecting a value corresponding to the entry.
10. The method of claim 1, wherein the statistical data comprises at least one of:
 - a count of accesses;
 - a count of page faults;
 - a time record; or
 - a generation counter.
11. The method of claim 1, wherein the application logic comprises at least one of:
 - a virtualization instance;
 - a virtual machine;
 - a container;
 - a jail; or
 - a zone.
12. The method of claim 1, wherein the invalidating comprises:
 - invalidating a page table entry associated with at least one virtual address in the first virtual address space at the client device; or
 - marking the page table entry associated with the at least one virtual address as not present.
13. The method of claim 1, further comprising:
 - selecting, by the client device, in response to a second request, a second portion of the local primary memory to reclaim; and
 - in response to the external primary memory being insufficient for the second request, identifying a portion of a swap space, and copying data from the second portion of local primary memory to the portion of the swap space.
14. The method of claim 13, wherein the swap space comprises at least one of:
 - a swap file;
 - a swap partition;
 - a remote swap device which is accessible via a first storage protocol; or

119

a local swap device which is accessible over a local bus via a second storage protocol.

15. The method of claim 1, wherein before copying the data from the first portion of local primary memory to the portion of external primary memory, the method further comprises:

marking at least one page table entry in the first virtual address space as one of: not writable; invalid; clean; not dirty; not modified; or not present.

16. The method of claim 1, wherein before copying the data from the first portion of local primary memory to the portion of external primary memory, the method further comprises:

unmapping at least one page table entry in the first virtual address space; or

restricting access to the at least one page table entry in the first virtual address space.

17. The method of claim 1, wherein the invalidating comprises converting the portion of the first virtual address space at the client device by remapping at least one virtual address in the first virtual address space at the client device from the first portion of local primary memory to the portion of external primary memory.

18. The method of claim 17, wherein at least a portion of a second virtual address space is mapped to the first portion of local primary memory in the client device before the first request, the method further comprising:

remapping the at least the portion of the second virtual address space to the portion of external primary memory.

19. The method of claim 17, further comprising:

in response to the first virtual address space being accessed or modified, aborting remapping the at least one virtual address in the first virtual address space.

120

20. A client device comprising:
a processor; and

a local primary memory in communication with the processor, the local primary memory comprising: an application logic unit; and a client logic unit configured to:

allocate, for the application logic unit, a region of external primary memory included in a memory appliance;

select, in response to a first request, a first portion of local primary memory in the client device to reclaim based on statistical data maintained for portions of the local primary memory;

determine a portion of external primary memory from the region of external primary memory;

copy data from the first portion of local primary memory to the portion of external primary memory; and

invalidate a portion of a first virtual address space at the client device corresponding to the first portion of the local primary memory.

21. A non-transitory storage medium for storing computer readable instructions, the computer readable instructions, when executed by a processor of a client device, causing the processor to:

allocate, for an application logic, a region of external primary memory included in a memory appliance;

select, in response to a first request, a first portion of local primary memory in the client device to reclaim based on statistical data maintained for portions of the local primary memory;

determine a portion of external primary memory from the region of external primary memory;

copy data from the first portion of local primary memory to the portion of external primary memory; and

invalidate a portion of a first virtual address space at the client device corresponding to the first portion of the local primary memory.

* * * * *