

Pamięć wirtualna


Rafał Łasocha


Basic functions of virtual mem.

- create virtual machine environment for every process
- demand-paging – granularity for process residence than an entire address space
- virtual-address aliasing (shared objects at different virtual addresses)
- protection aliasing (shared resources may have different protections dependent on process)
- support for virtual caches

Memory is a cache for permanent store

Organizations of Physical Memory


Virtual Page 5



defs

- page tables – tables which keep mappings between virtual and physical memory
- page table entries – entries of former


What info OS wants to have from PTE

- virtual page number (VPN)
- physical frame number (PFN) or location on disk
- ID of the page's owner (ASID)
- page's protection information
- aid in making replacement decision (whether it was recently accessed / written)
- is it valid (valid bit)

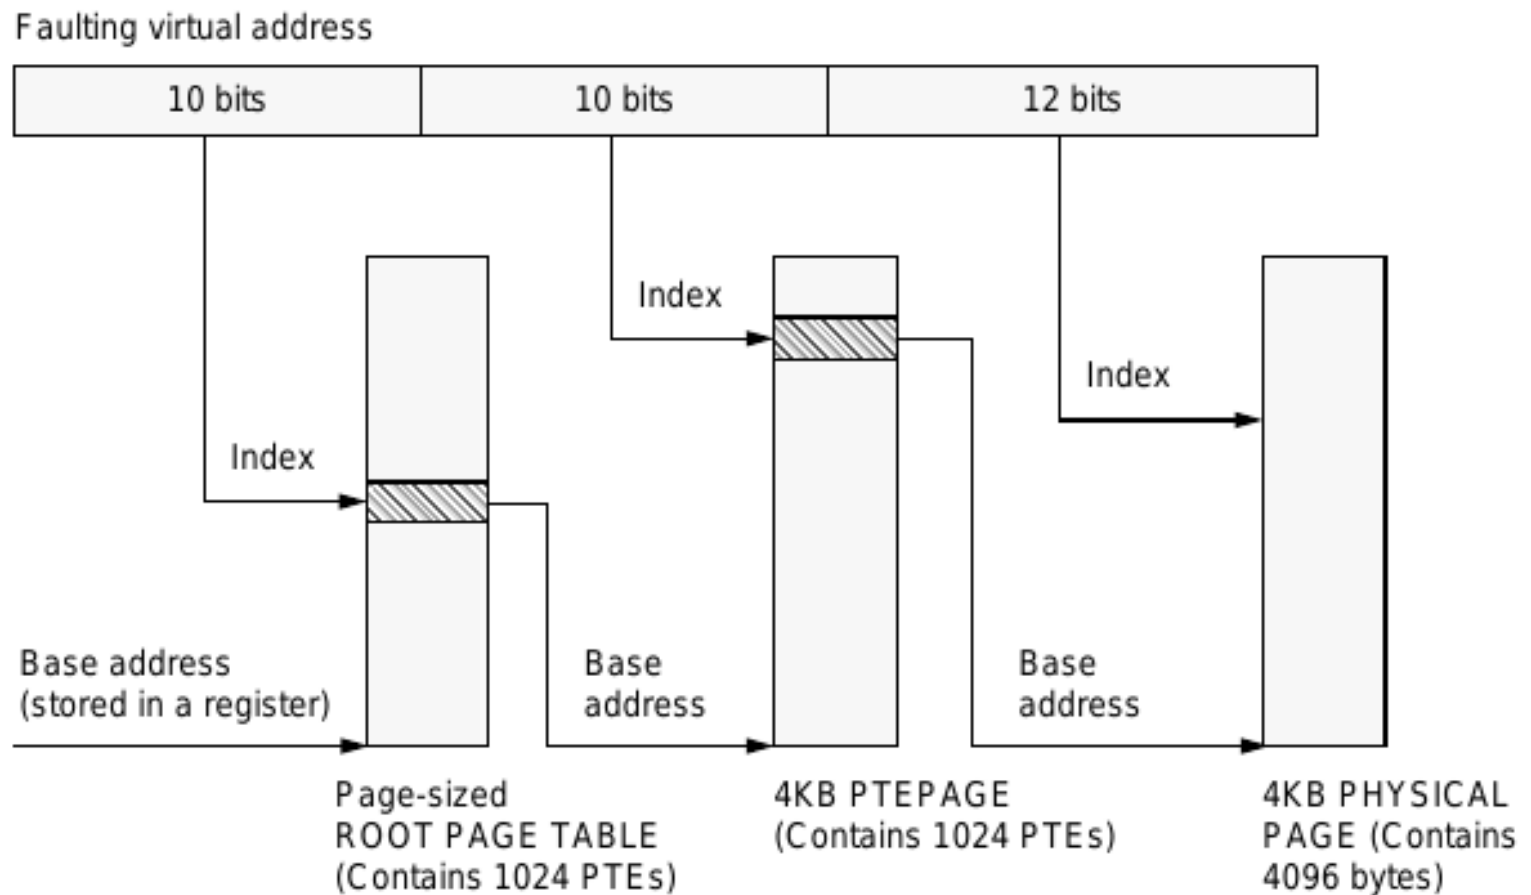
Page table organization

- hierarchical page table
 - top-down traversal (IA-32)
 - bottom-up traversal (MIPS, Alpha)
- inverted page table
 - basic idea
 - modifications in PA-RISC
 - modifications in PowerPC

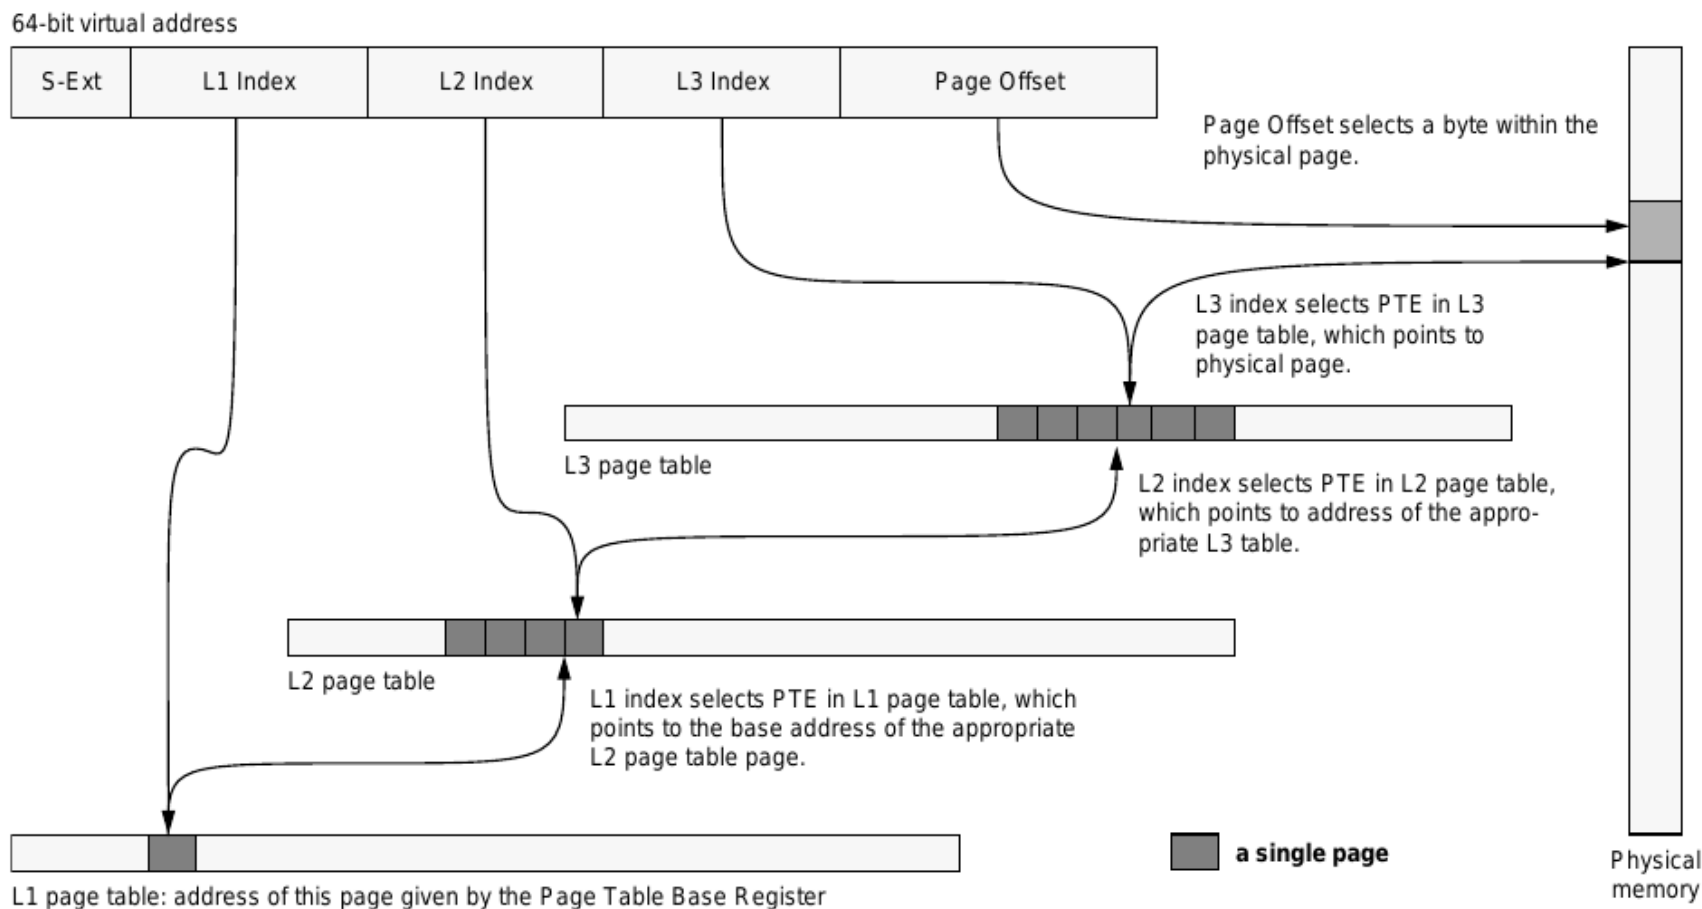
Hierarchical Page Table



Top-down approach



3-level page table in Alpha



Bottom-up approach

- 3 memory references to get data in top-down approach for 32-bit architecture
- even more for 64-bit
- User Page Table is contiguous
- we can use that fact and decrease memory references to 2 in best-case scenario (in both 32 and 64-bit architectures)
- and use top-down approach only when we have page fault

Bottom-up approach

Faulting virtual address:



Virtual address of PTE in User Page Table (UPTE):



Physical address of PTE in Root Page Table (RPTE):




1

2

Inverted page table

- instead of table indexed by VPN and entries containing PFN, let's index it with PFN, and keep entries to VPN (and ASID)
- PFN is implicit
- instead of scaling with virtual address space, we scale with physical memory
- we still want use that page table to find *PFN*, given *VPN*
 - linear lookup - nope
 - hashing to the rescue

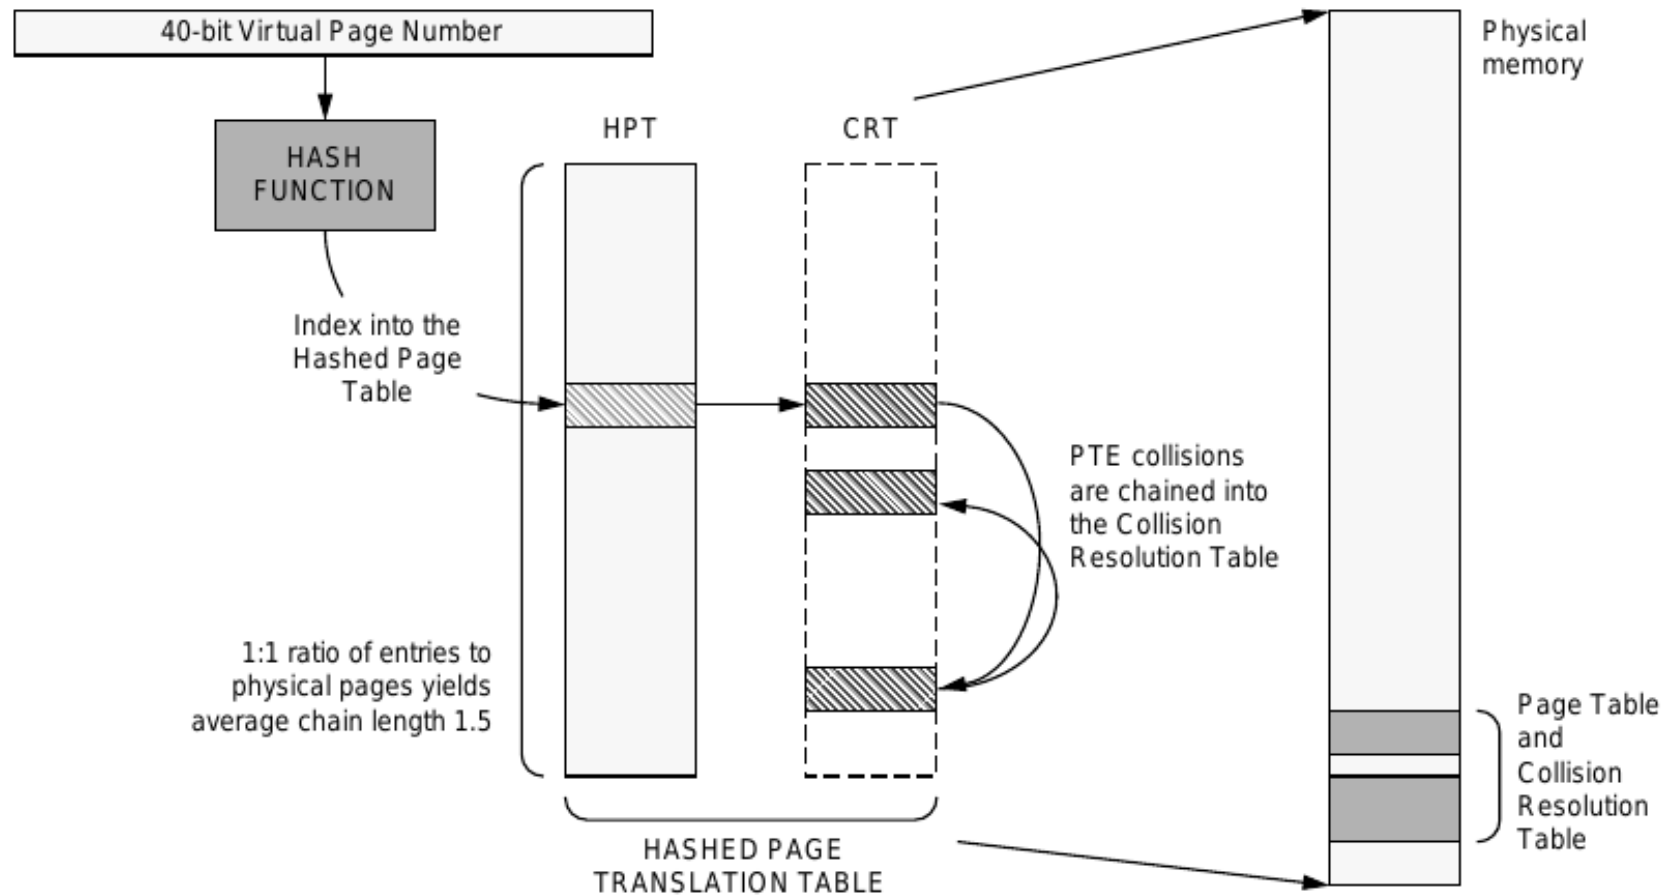
Inverted Page Table



PA-RISC

- at least 2 mem. ref. - we can do better than that
- let's remove access to HAT
- and just make IPT bigger
- disadvantage: we lose implicit information about PFN, so we have to keep it explicitly


PA-RISC



PowerPC

- let's limit chain's length by constant (8)
- and keep whole chain in contiguous space
- then we can fetch all PTE together and check whole chain at once
- however, if chain's length is bigger than 8, we delegate this problem to OS
- each PTE is also bigger, but we have only one (big) memory reference

PowerPC



Inverted vs. Hierarchical

- another structure is needed to keep references to disk
- inverted cannot into shared memory
- bigger PTE make communication with cache harder
- process sparsely used memory is wasting it

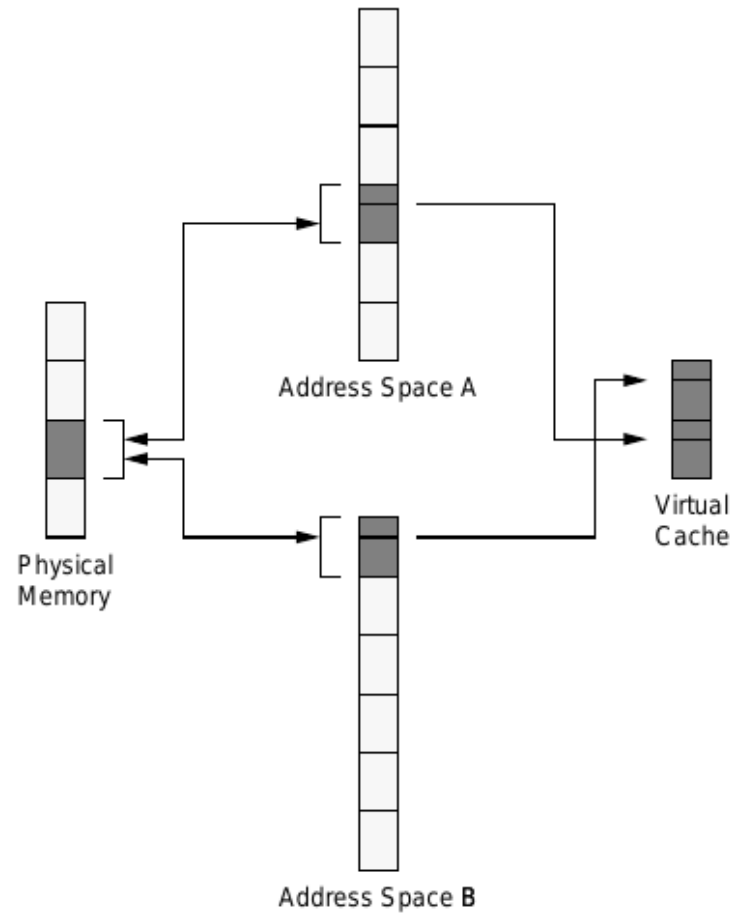
Translation lookaside buffer (TLB)

- page table cache (usually on CPU)
- some time ago, when address space were much smaller, one table mapping entire address space was small enough to be in hardware
- protection may block cache
- doesn't scale
- management: hardware, software (in OS)

hardware vs. software TLB

- state machine walking page table when TLB miss
- PowerPC, x86
- PA-7200 – hybrid approach, part of HPT in hardware
- slower
- interrupts
- flexible
- MIPS, SPARC, Alpha, PA-RISC

Synonym problem




Кэш-память

Процесс А



Процесс Б



Индекс

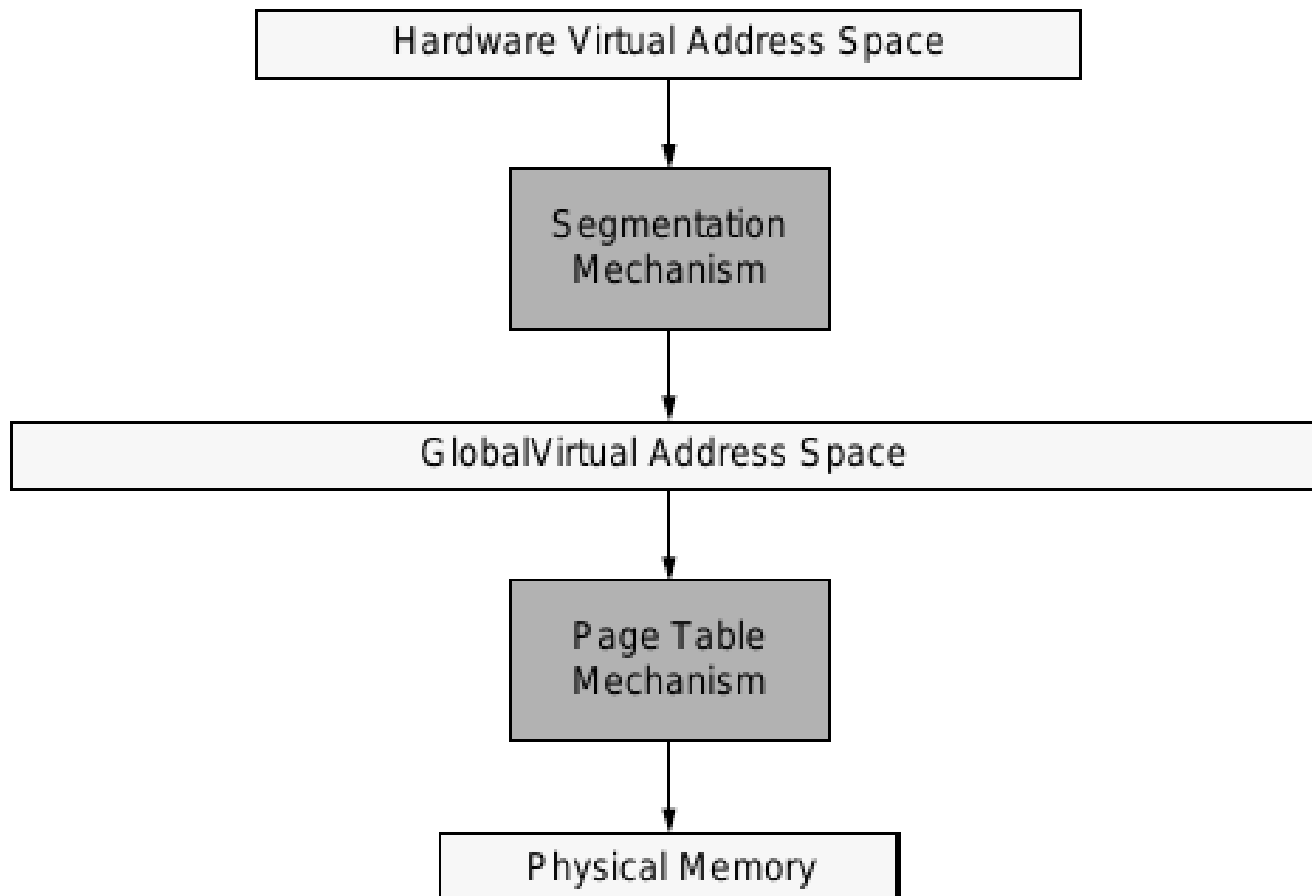
- 0
- 1
- 2
- 3
- 4
- 5
- 6
- 7

Канал 1 (Way 1)

Канал 2 (Way 2)

Канал 1 (Way 1)		Канал 2 (Way 2)	
Тэг	Строка (64 байта)	Тэг	Строка (64 байта)
		0x0	Строка 0
0x0	Строка 1		
0x0	Строка 2		
		0x0	Строка 3
0x0	Строка 4		

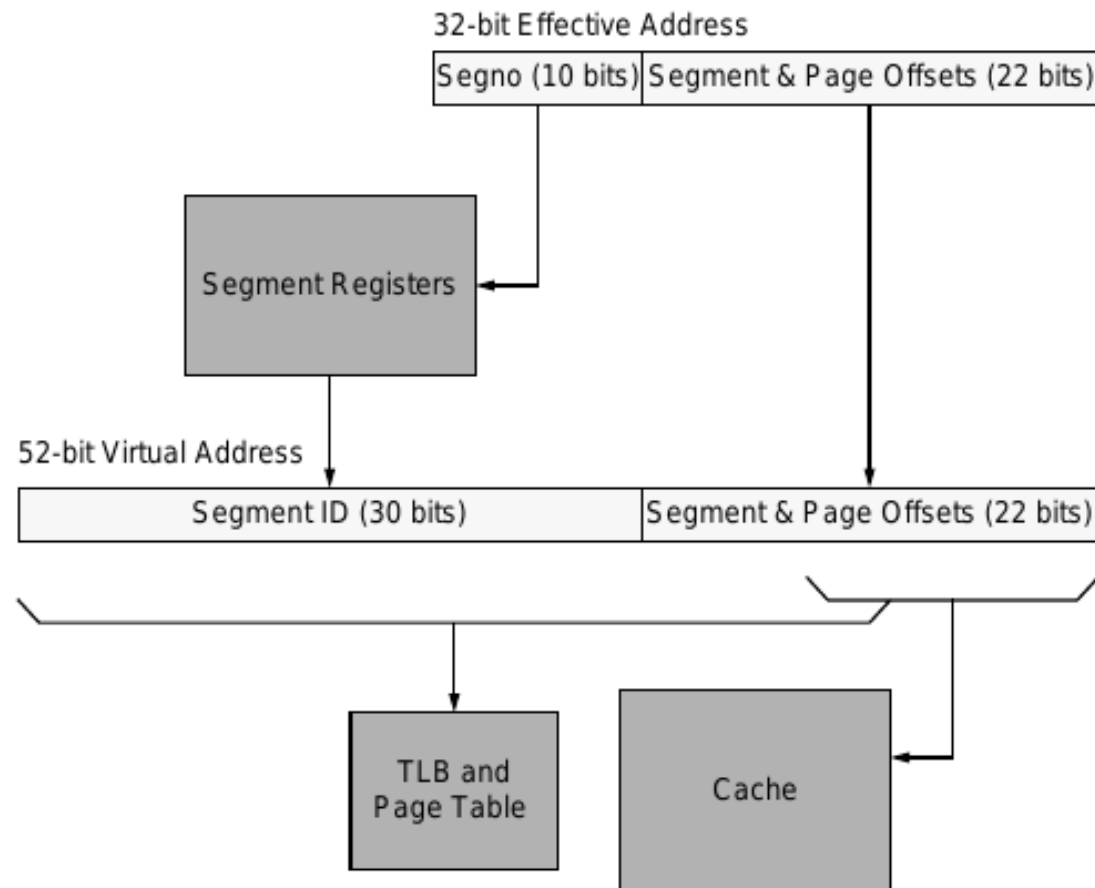
Global segmented space



Global segmented space

- if there's one global space, synonym problem disappears immediately
- segments are groups of pages
- segments may be scattered through whole global space
- automatic address-space protection

Global segmented space (example)



Global segmented space (example)

- per-user RPT - 4KB, UPT - 4MB, User Address Space - 4GB
- whole space: $2^{52} = 4\text{PB}$
- Global Page Table - contiguous, 4TB in the end of whole global space
- per-user RPT are small enough (just like RPT) to be wired down in physical memory

Global segmented space (example)

Segment Table: Segment Identifiers
(1024 30-bit IDs) plus protection bits



Per-User Root Page Table
4KB: 1024 PTEs



A 4-byte PTE,
which maps 4KB

4 B

Maps

A 4KB PTE Page: a continuous group
of 1024 PTEs that collectively map 4MB

4 KB

Maps

A 4MB virtual segment,
1/1024 of an address space

4 MB


4 MB

Virtual Segment

Unmapped Physical Memory


Mapped Virtual Memory

User Page Table
4MB: 1024 4KB PTE Pages, 1024K PTEs




Pages are not contiguous in virtual space;
they are found at the top of the global space

Per-Process User Address Space
4GB: 1024 4MB segments, 1024K 4KB virtual pages



52-bit Global Virtual Address Space: 2^{30} 4MB virtual segments

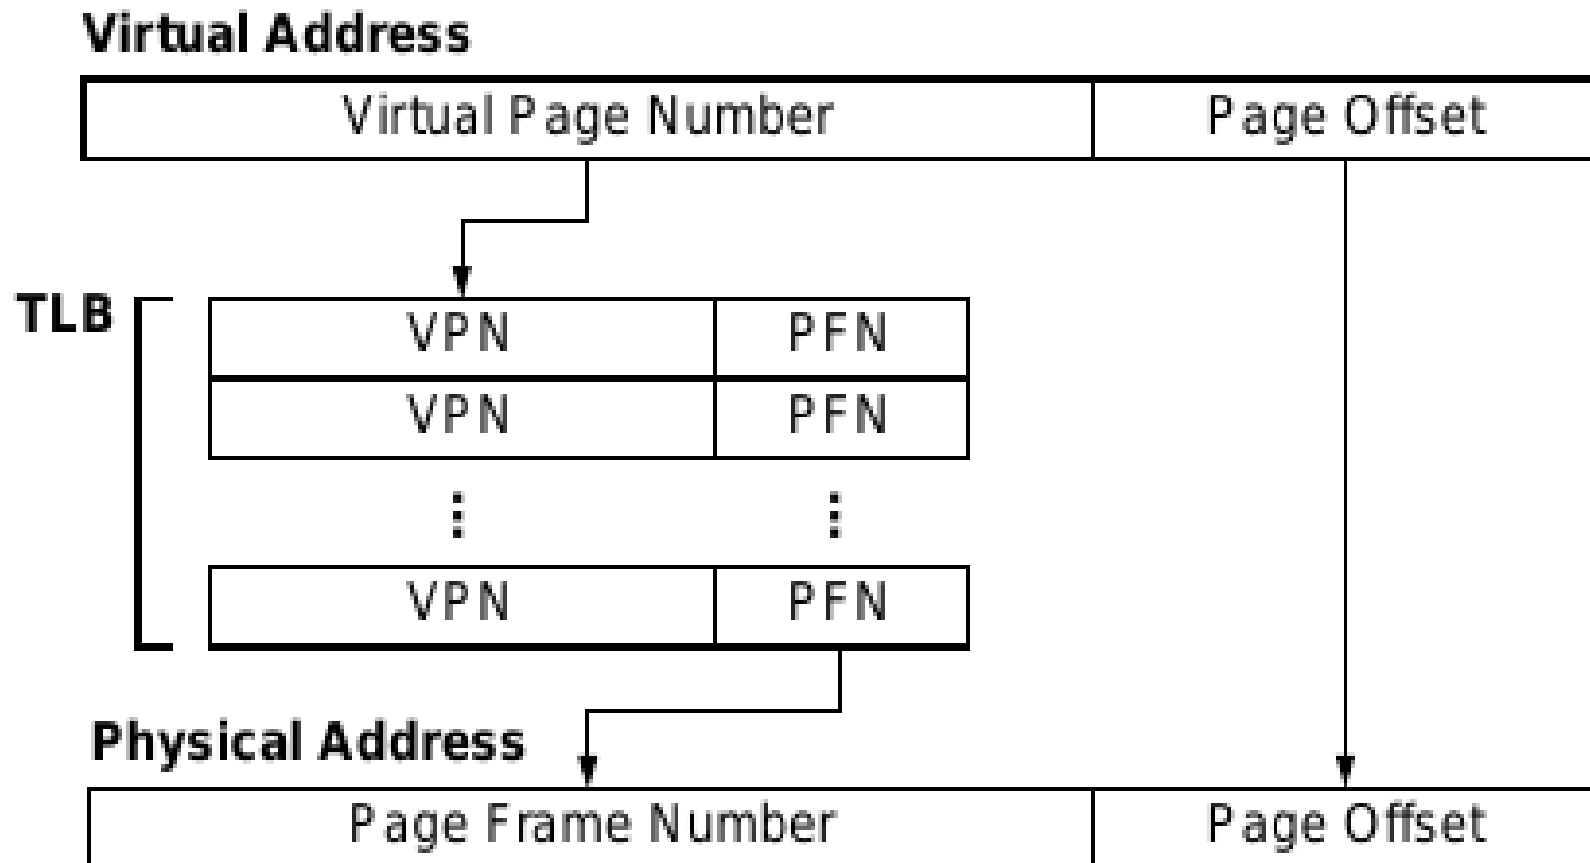


2^{30} PTE Pages: 4KB each,
4TB total

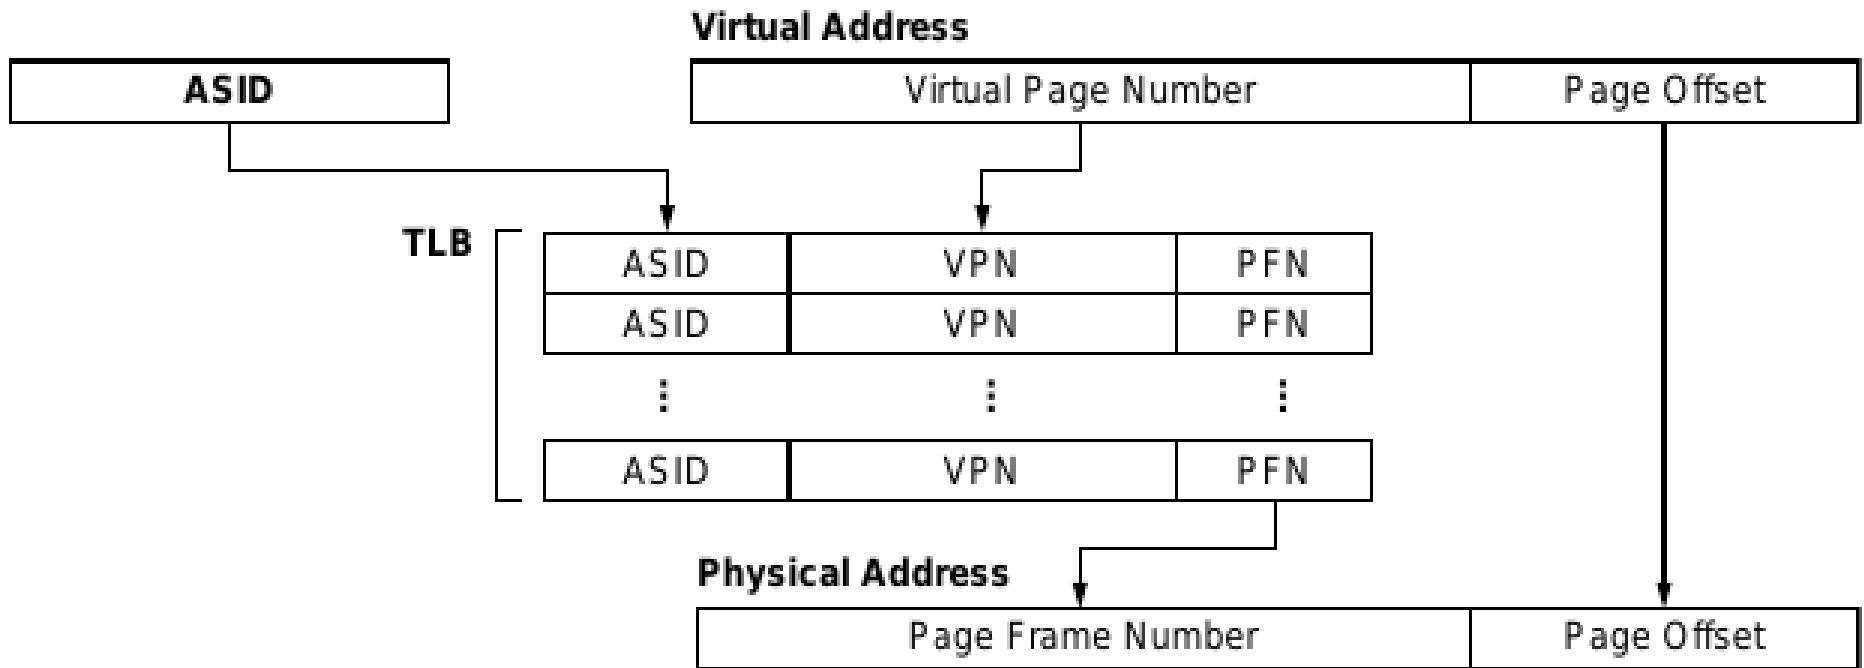
Address Space Organizations

- virtual space may be categorized by:
 - owner:
 - single owner (basic implementation)
 - multiple owner (segmentation)
 - used identifiers (protection)
 - no ID
 - single ID
 - multiple ID

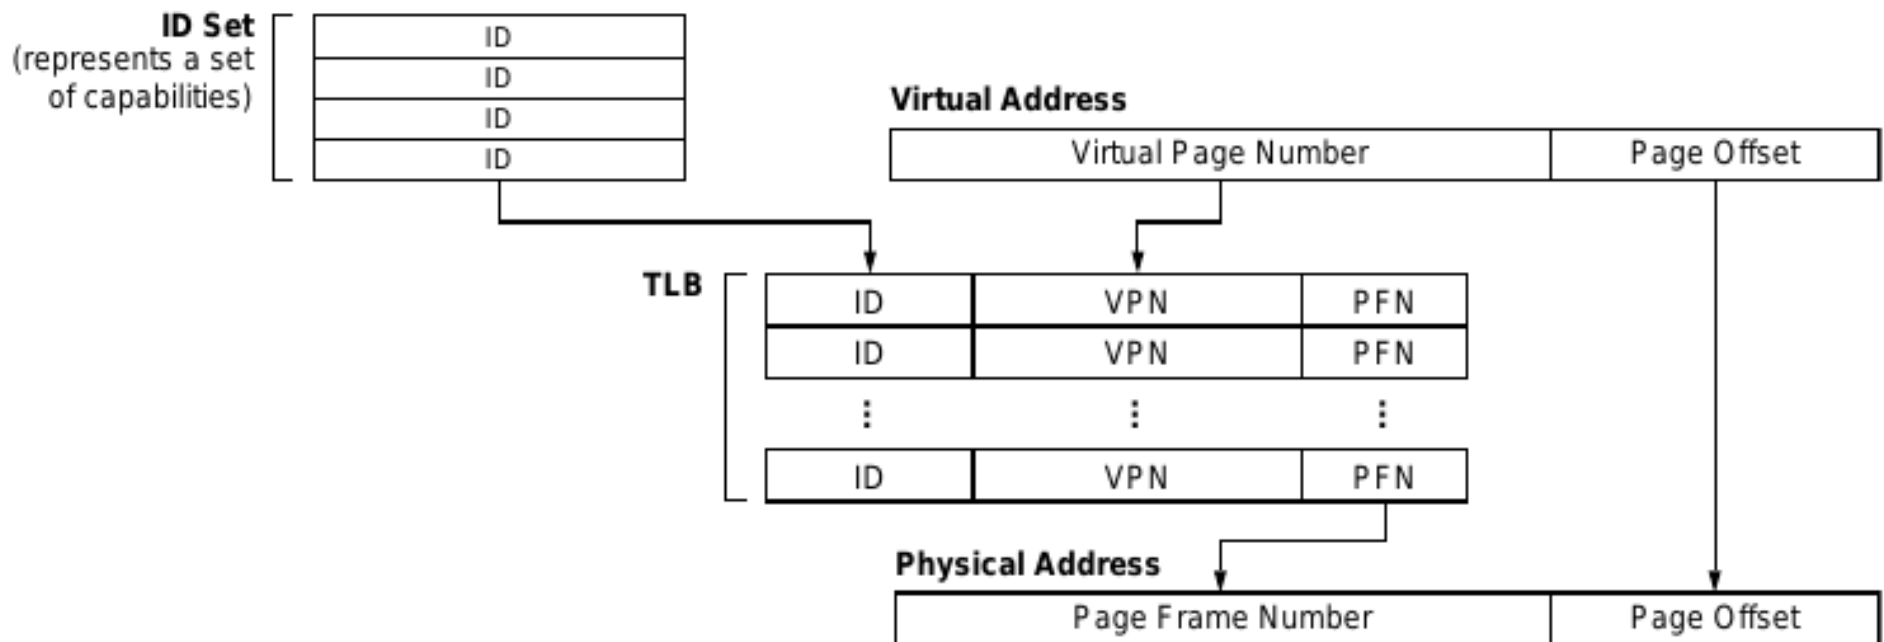
Single owner, no ID



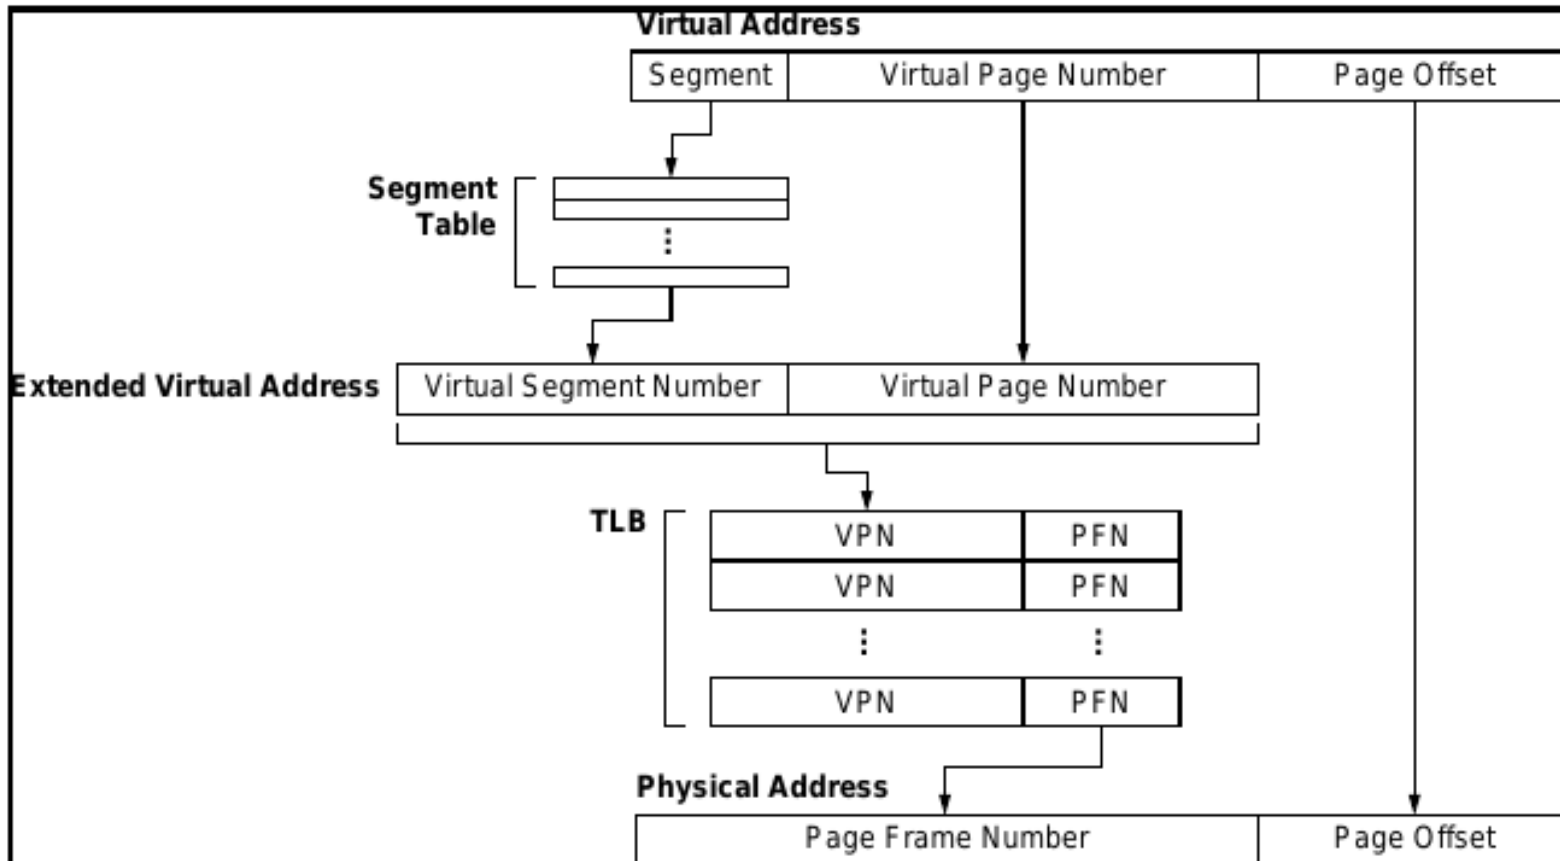
Single owner, single ID



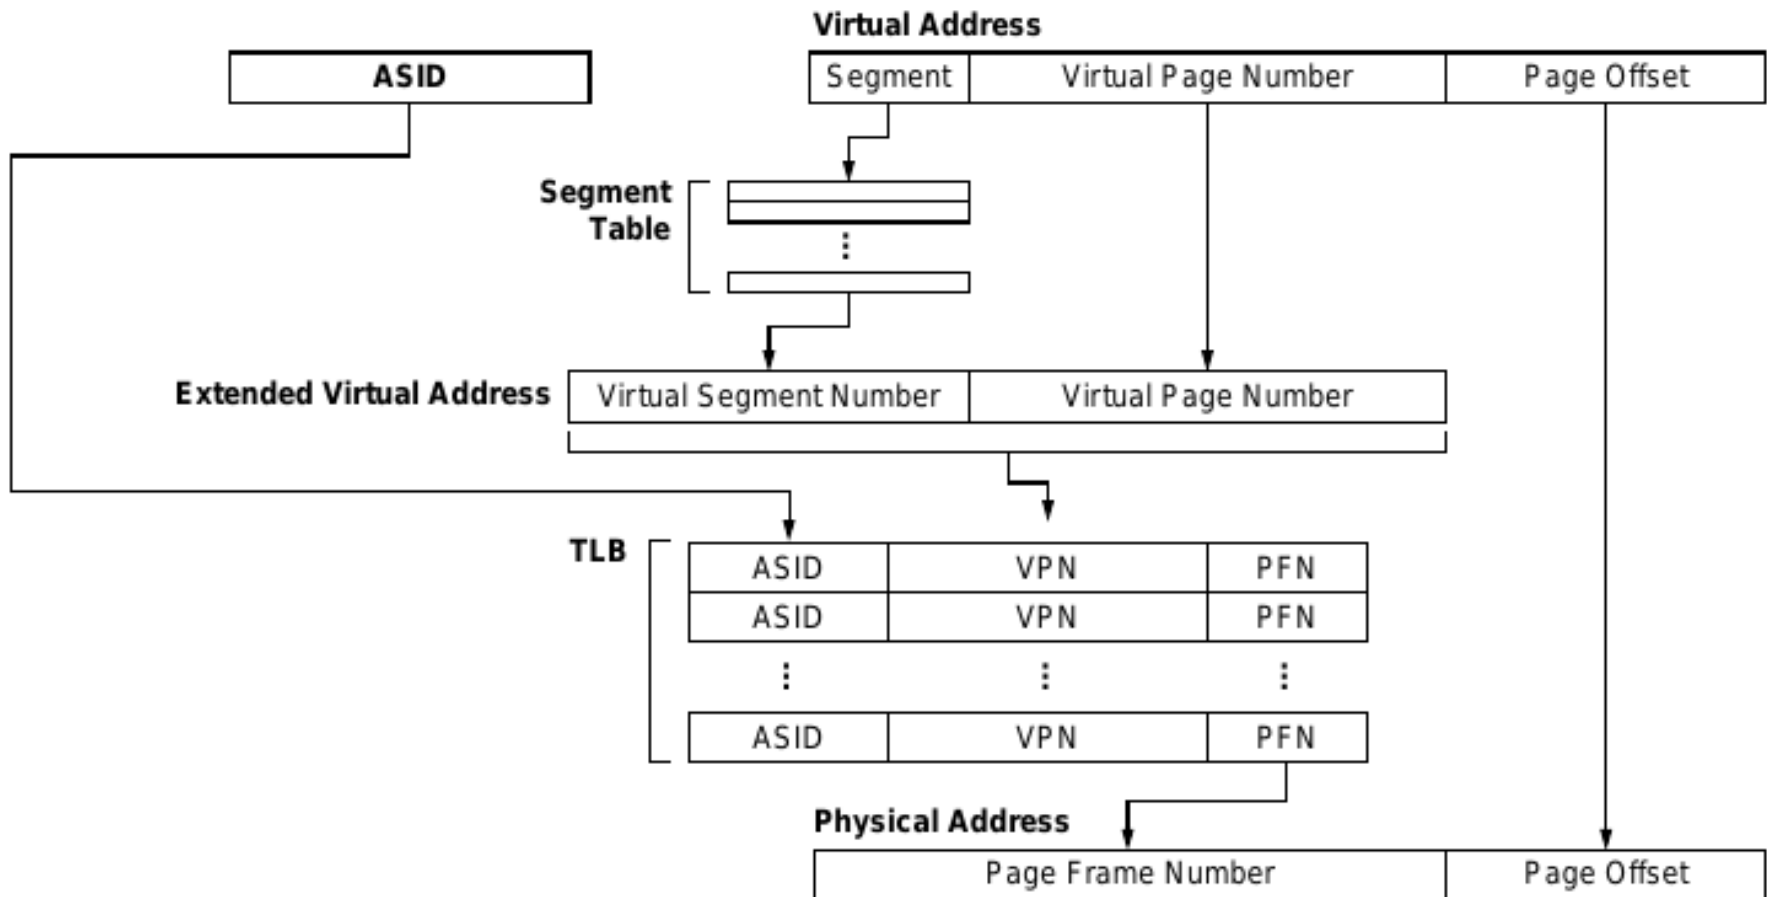
Single owner, multiple ID



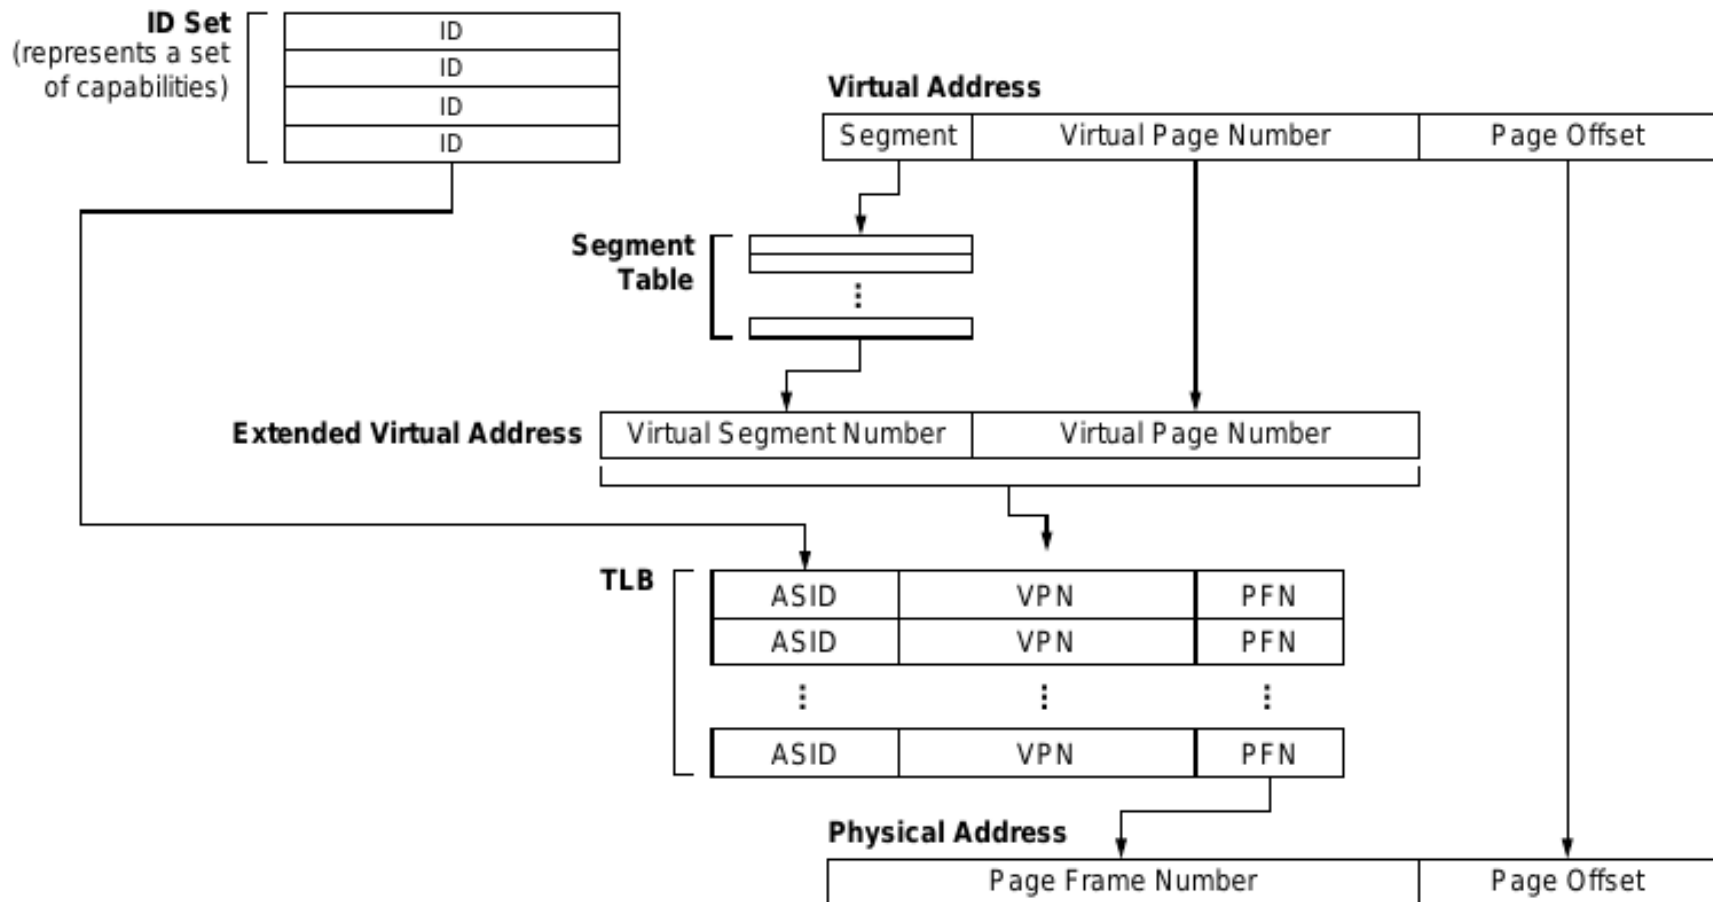
Multiple owner, no id

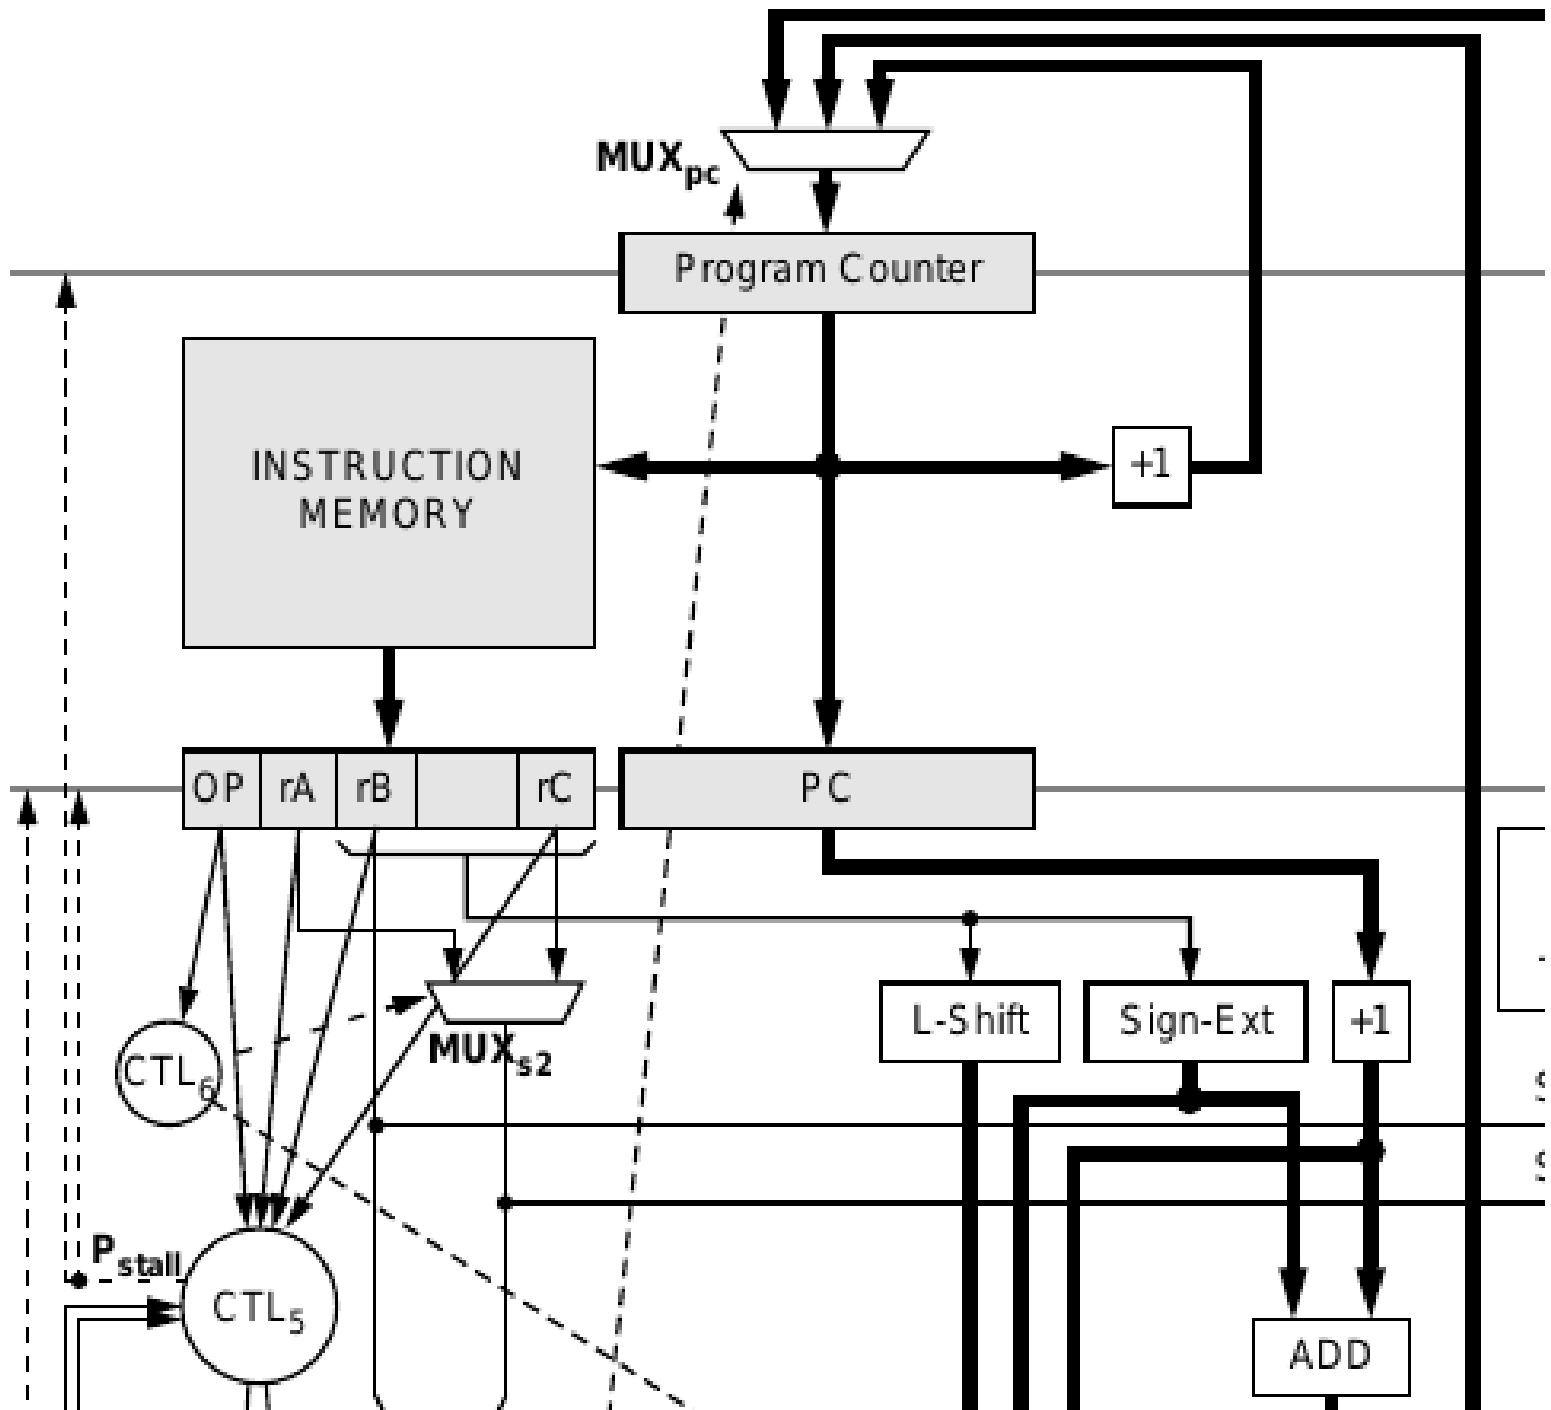


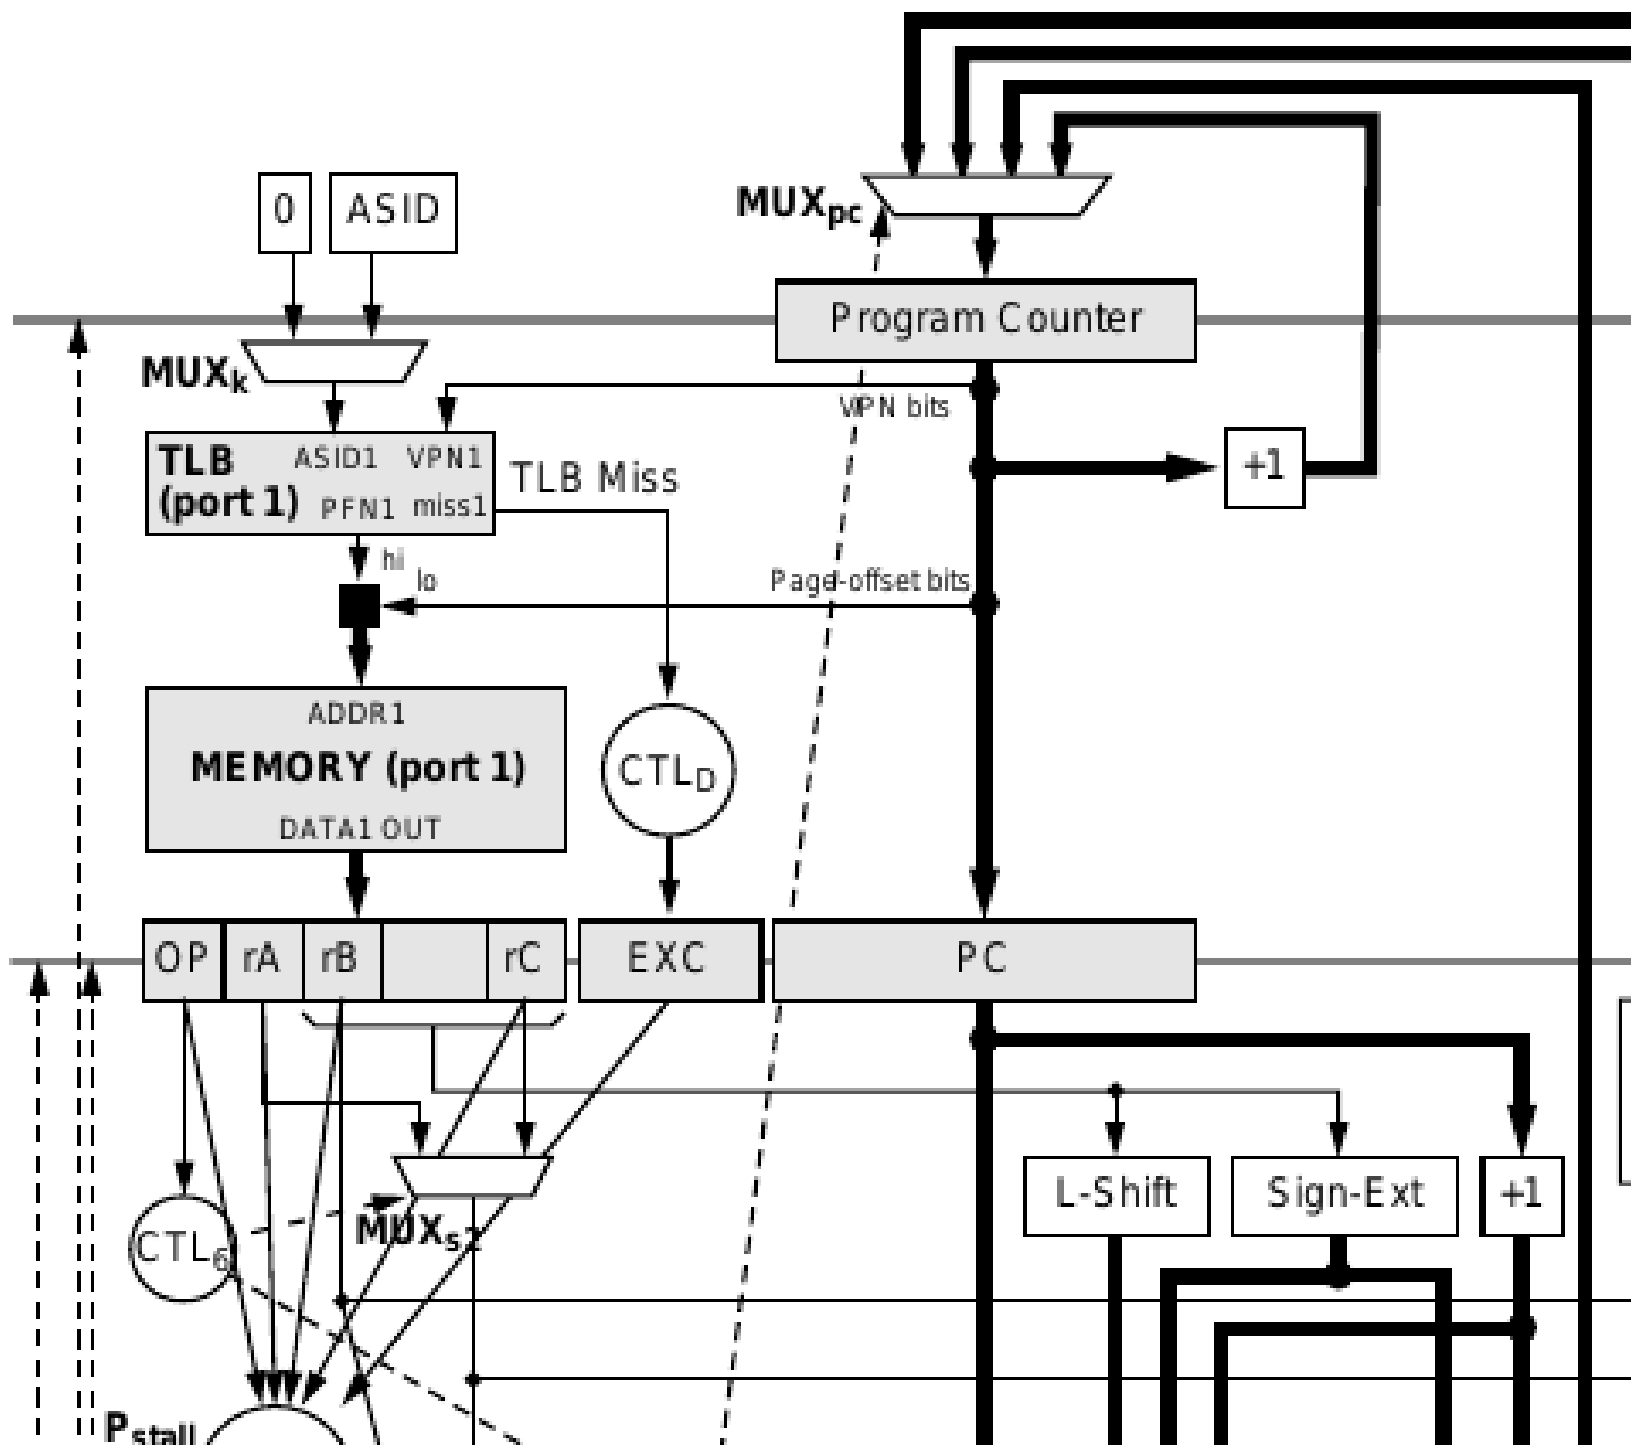
Multiple owner, single ID

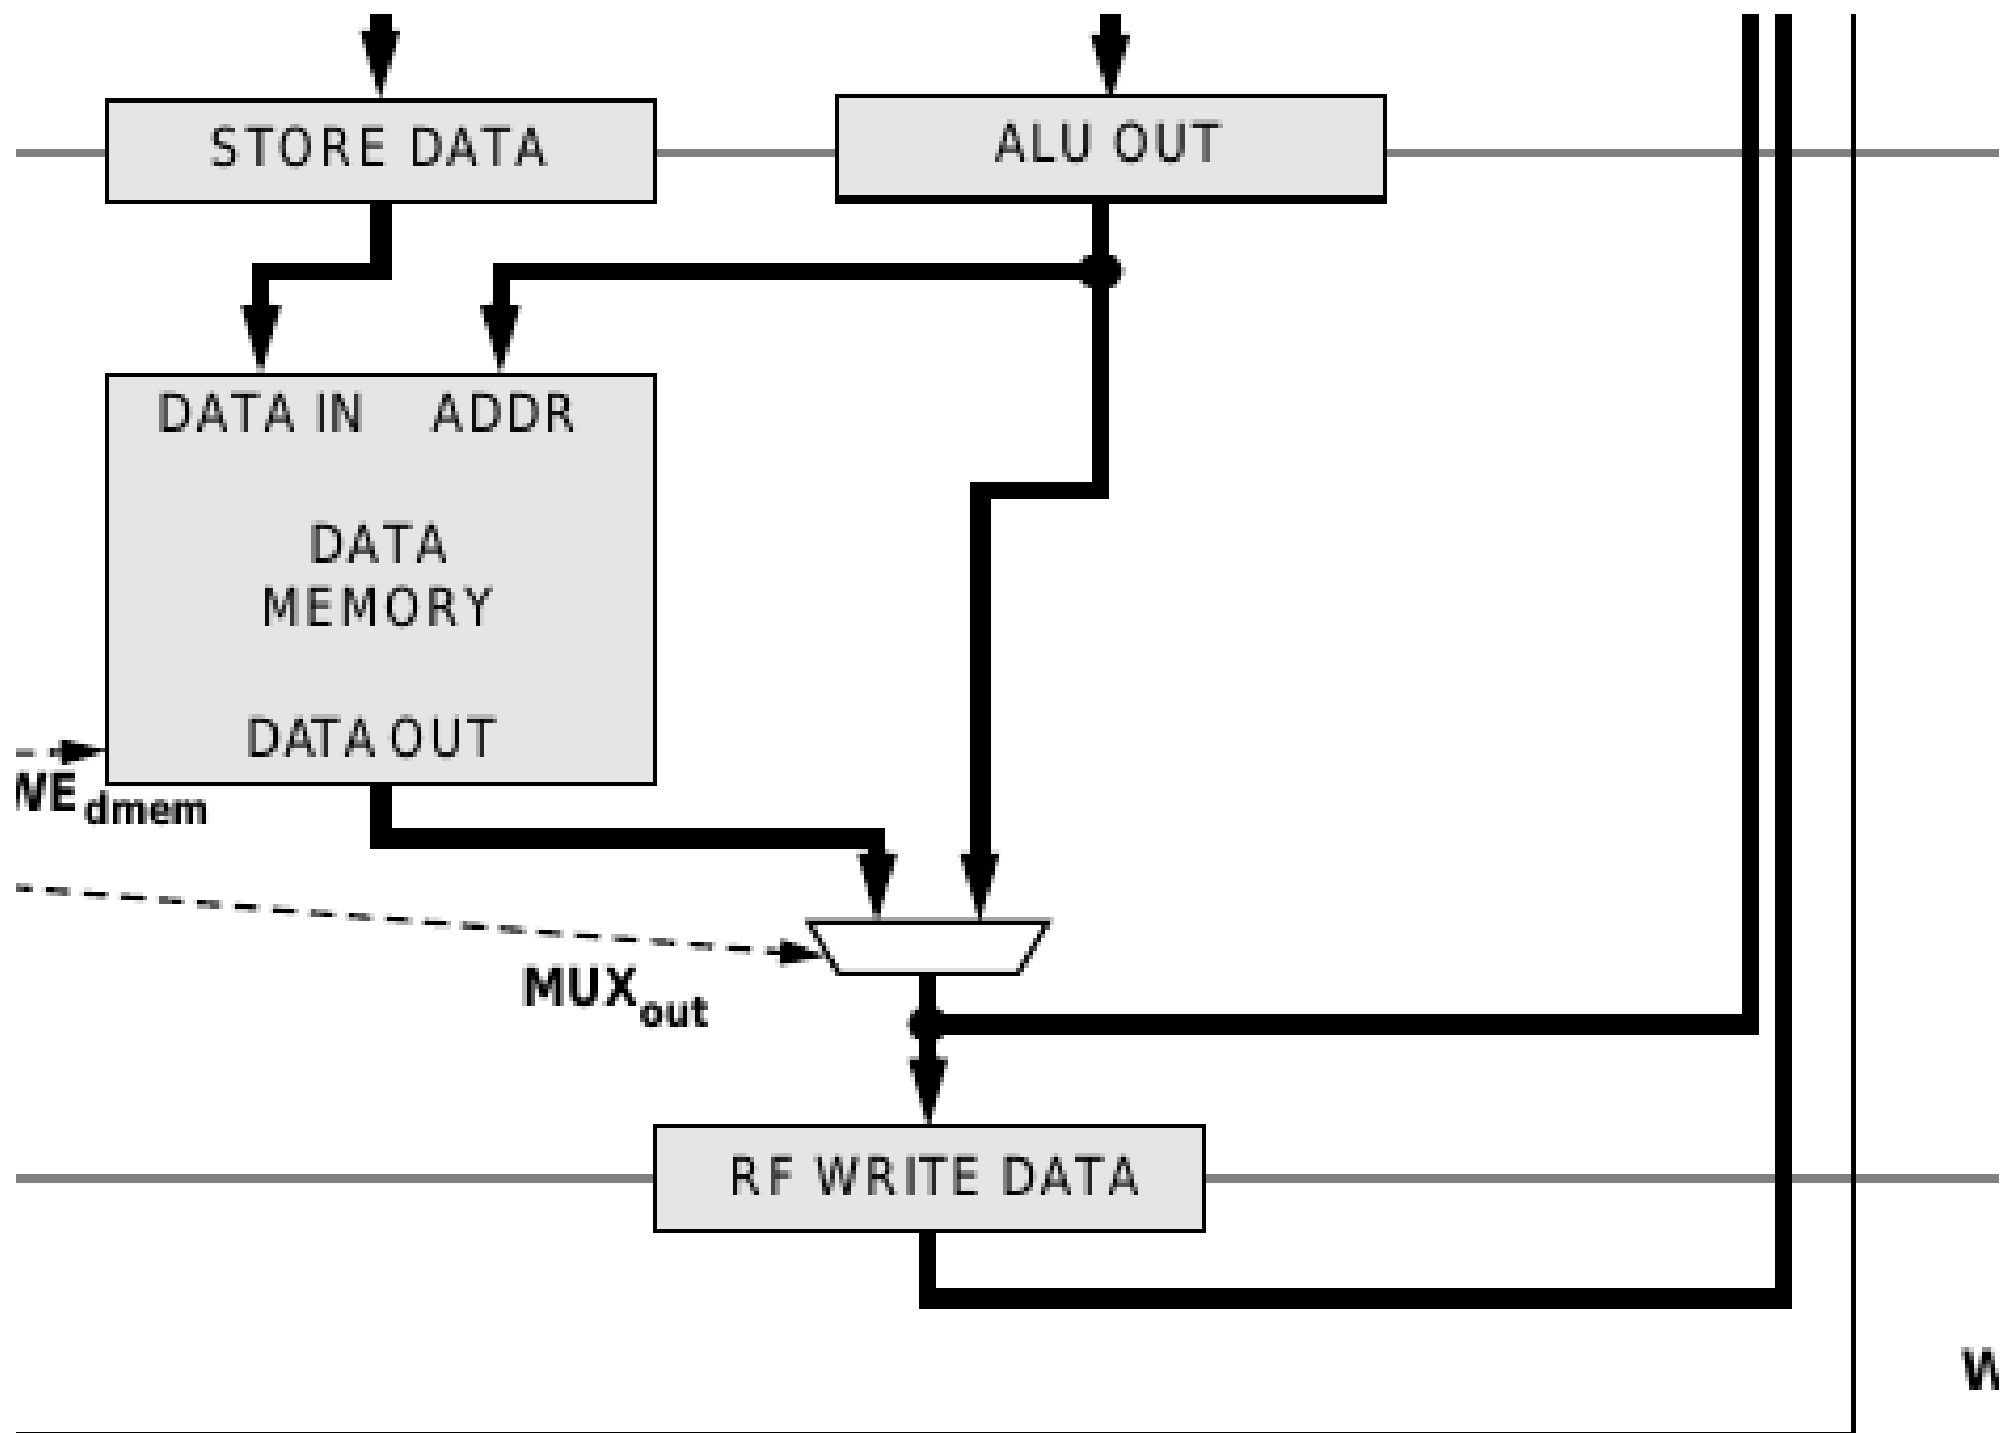


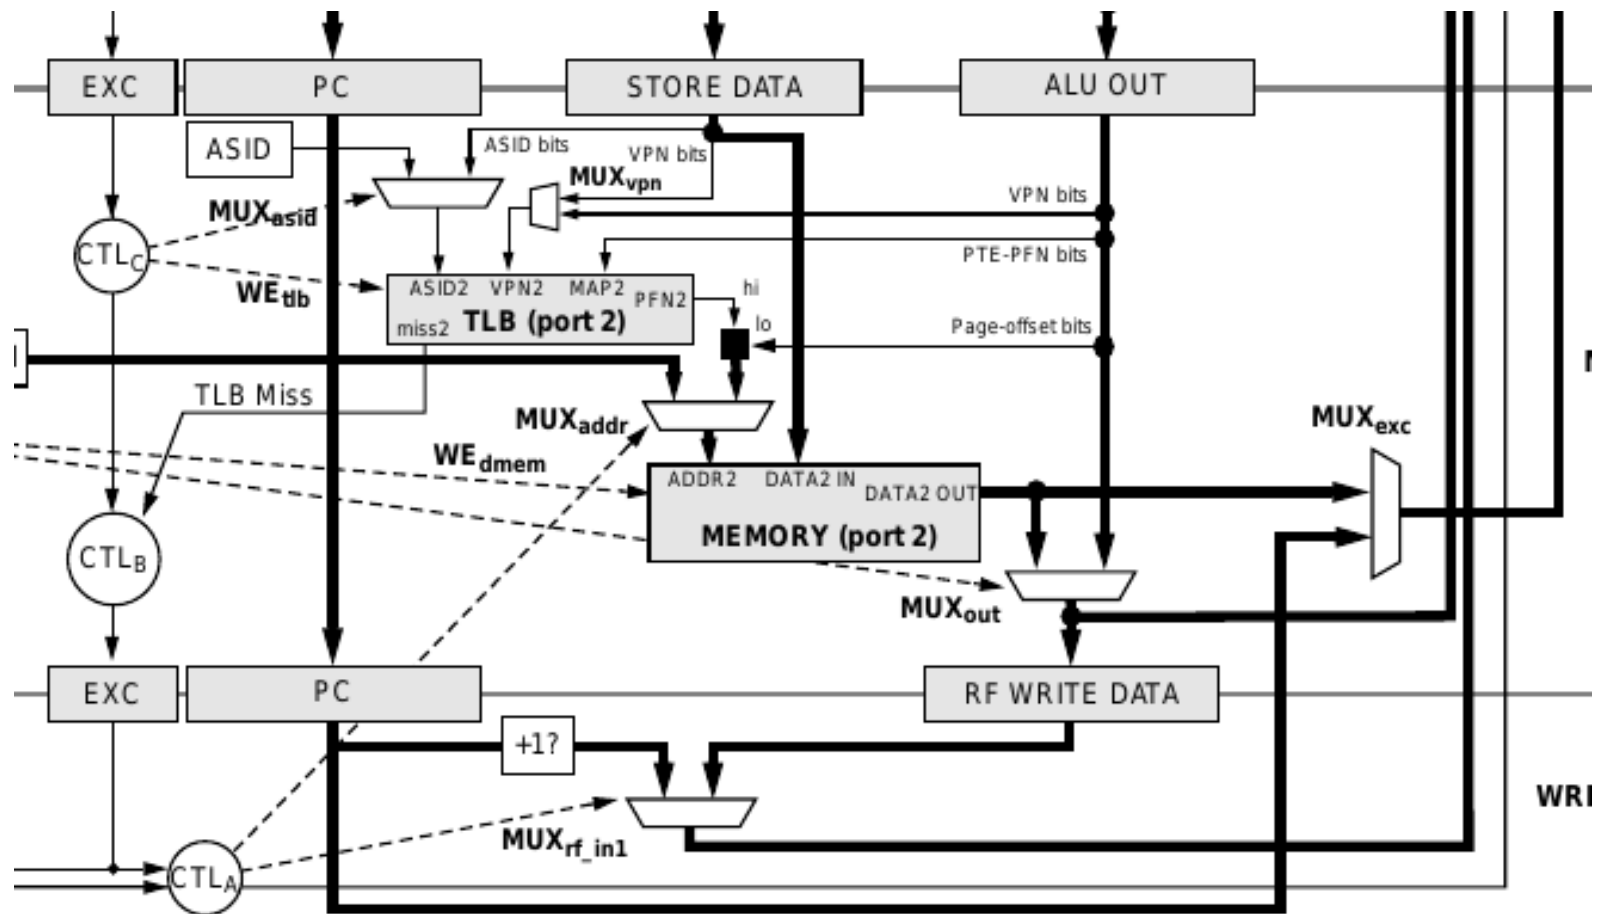
Multiple owner, multiple ID











questions?