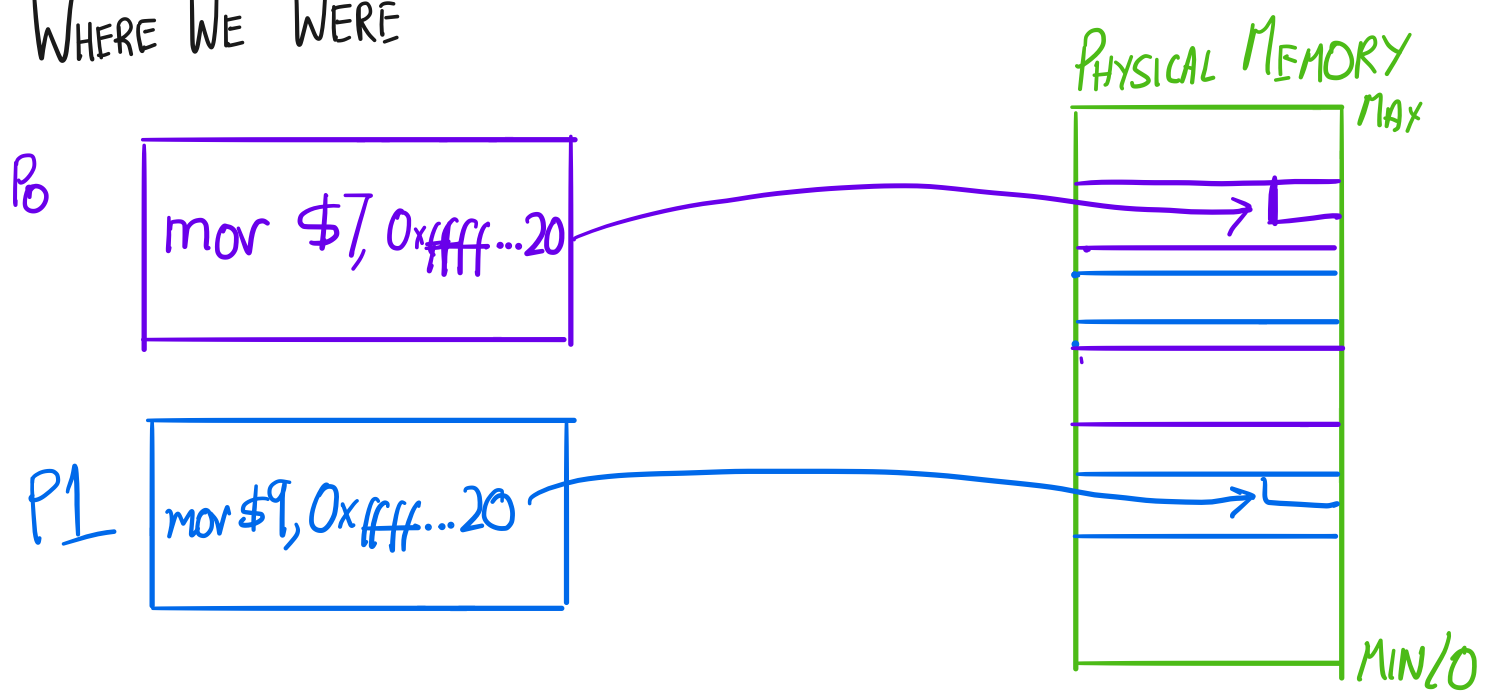
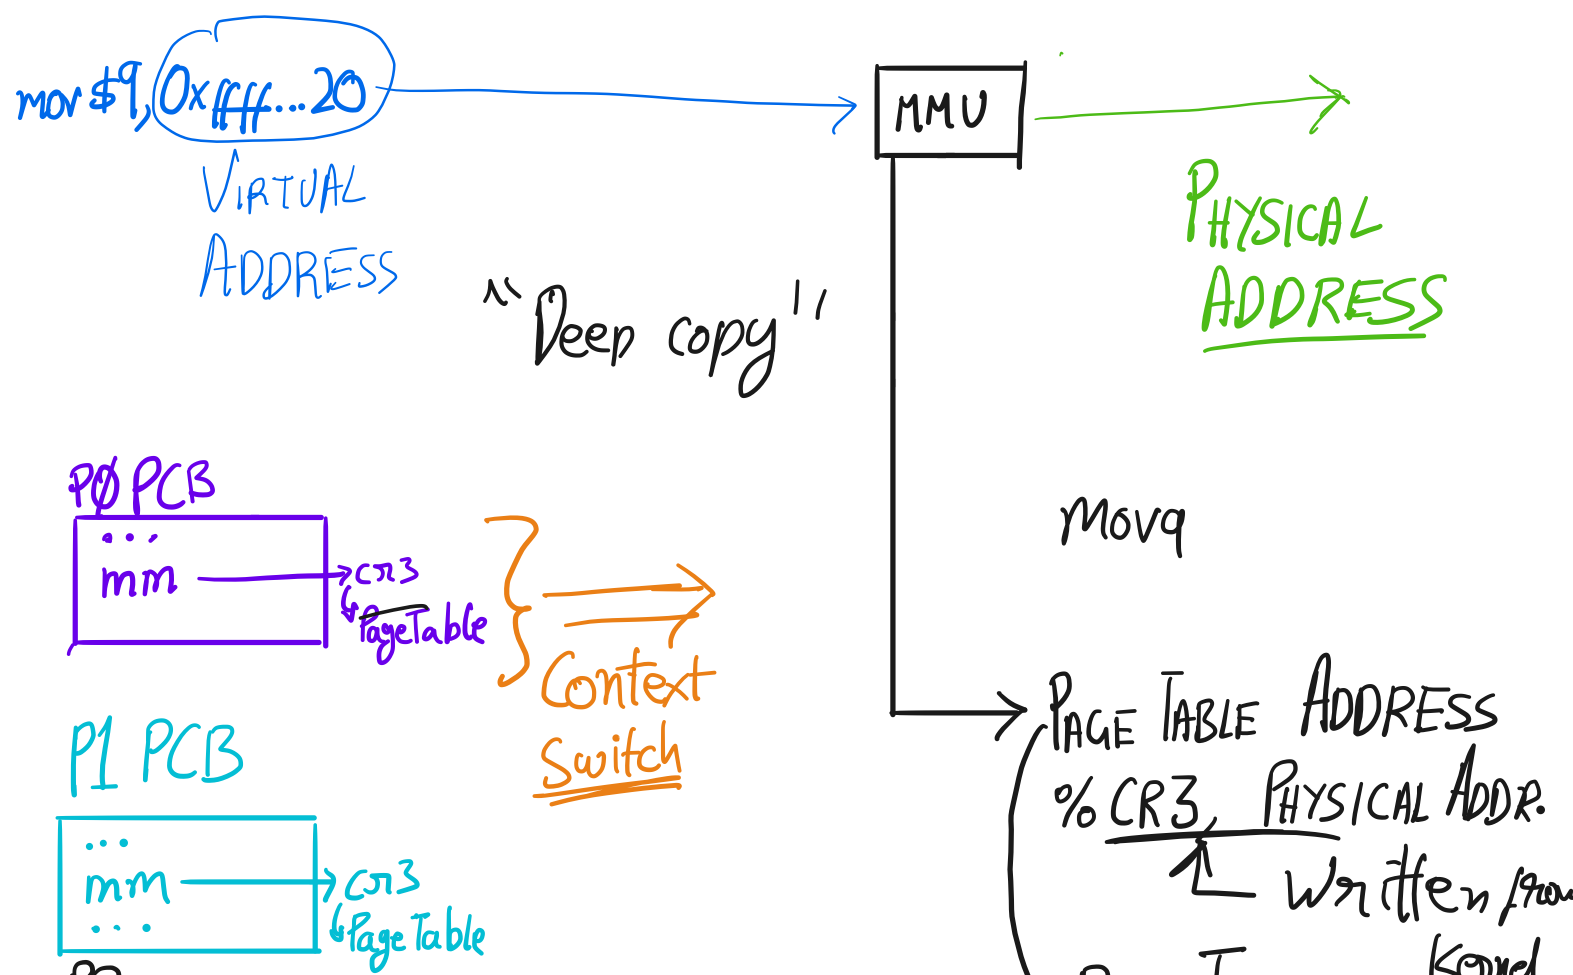


# CS202 - VIRTUAL MEMORY

WHERE WE WERE



How?

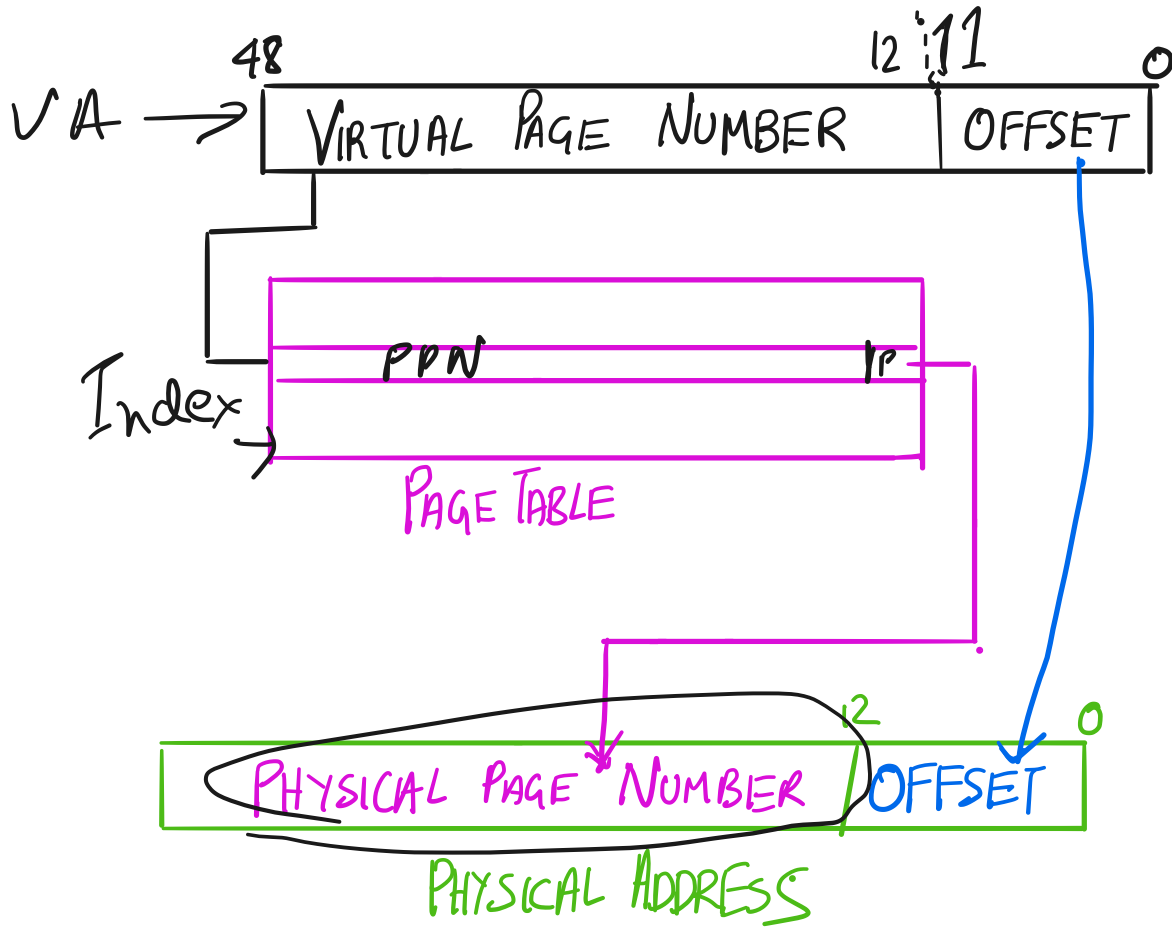


mm

↓ PAGE TABLES

# TODAY: PAGE TABLES

VIRTUAL ADDRESS (FOR NOW ASSUME 48-BITS)



NOTE PHYSICAL & VIRTUAL ADDRESSES MIGHT HAVE DIFFERENT LENGTHS

AMD64 (some) :- Virtual 48 bits  
Physical 52 bits

NOTE REMAINING  $64-48 \equiv 16$  bits?

ALL 0 OR ALL 1

## PAGE TABLES

CORE PROBLEM: LOTS OF UNUSED VIRTUAL ADDRESSES

$2^{36} \equiv$

DON'T WANT TO WASTE SPACE.

SOLUTION: BUILD A TREE

(SWITCH TO SHEET)

## WHAT IS IMPORTANT


- FAULTS: NOT PRESENT OR  
VIOLATE ACCESS CONTROL PERMISSIONS

... ADDRESS

- How to TRANSLATE VIRTUAL ADDRESSES
  - VPN, OFFSET
  - USE VPN TO WALK PAGE TABLE
- WHY THE TREE HELPS.

uint64\_t x[512]

## RESOLVING ADDRESSES



- RED
- 4 MEMORY ACCESSES TO LOCATE PPN
- + 1 Actual access

- 5x EXTRA ACCESSES

- WANT TO REDUCE OVERHEADS.

→ CACHE RECENT VPN → PPN

TRANSLATION  
LOOKASIDE  
BUFFER

TRANSLATIONS

→ SIZE: DICTATED BY HARDWARE

USUALLY B/W 64-1024 entries

- ON MISS?

TLB QUESTION

16 ENTRIES.

PROGRAM THAT (AFTER INITIAL START) MISSES TLB ON  
EACH INSTRUCTION? Q [4096 \* 12]

~~for~~ (int i=0; i<4096; i++)  
while (i)

puts(a[i])

i = i + 4096

# PROCESS MEMORY LAYOUT.

Q096. (7)

$2^0$

Byte

$2^{10}$

Kilobyte

(1024)

$2^{20}$

MEGA BYTE

$2^{30}$

GIGA BYTE

...

$2^3, 2^9$

< 1 PT

↳ Entries

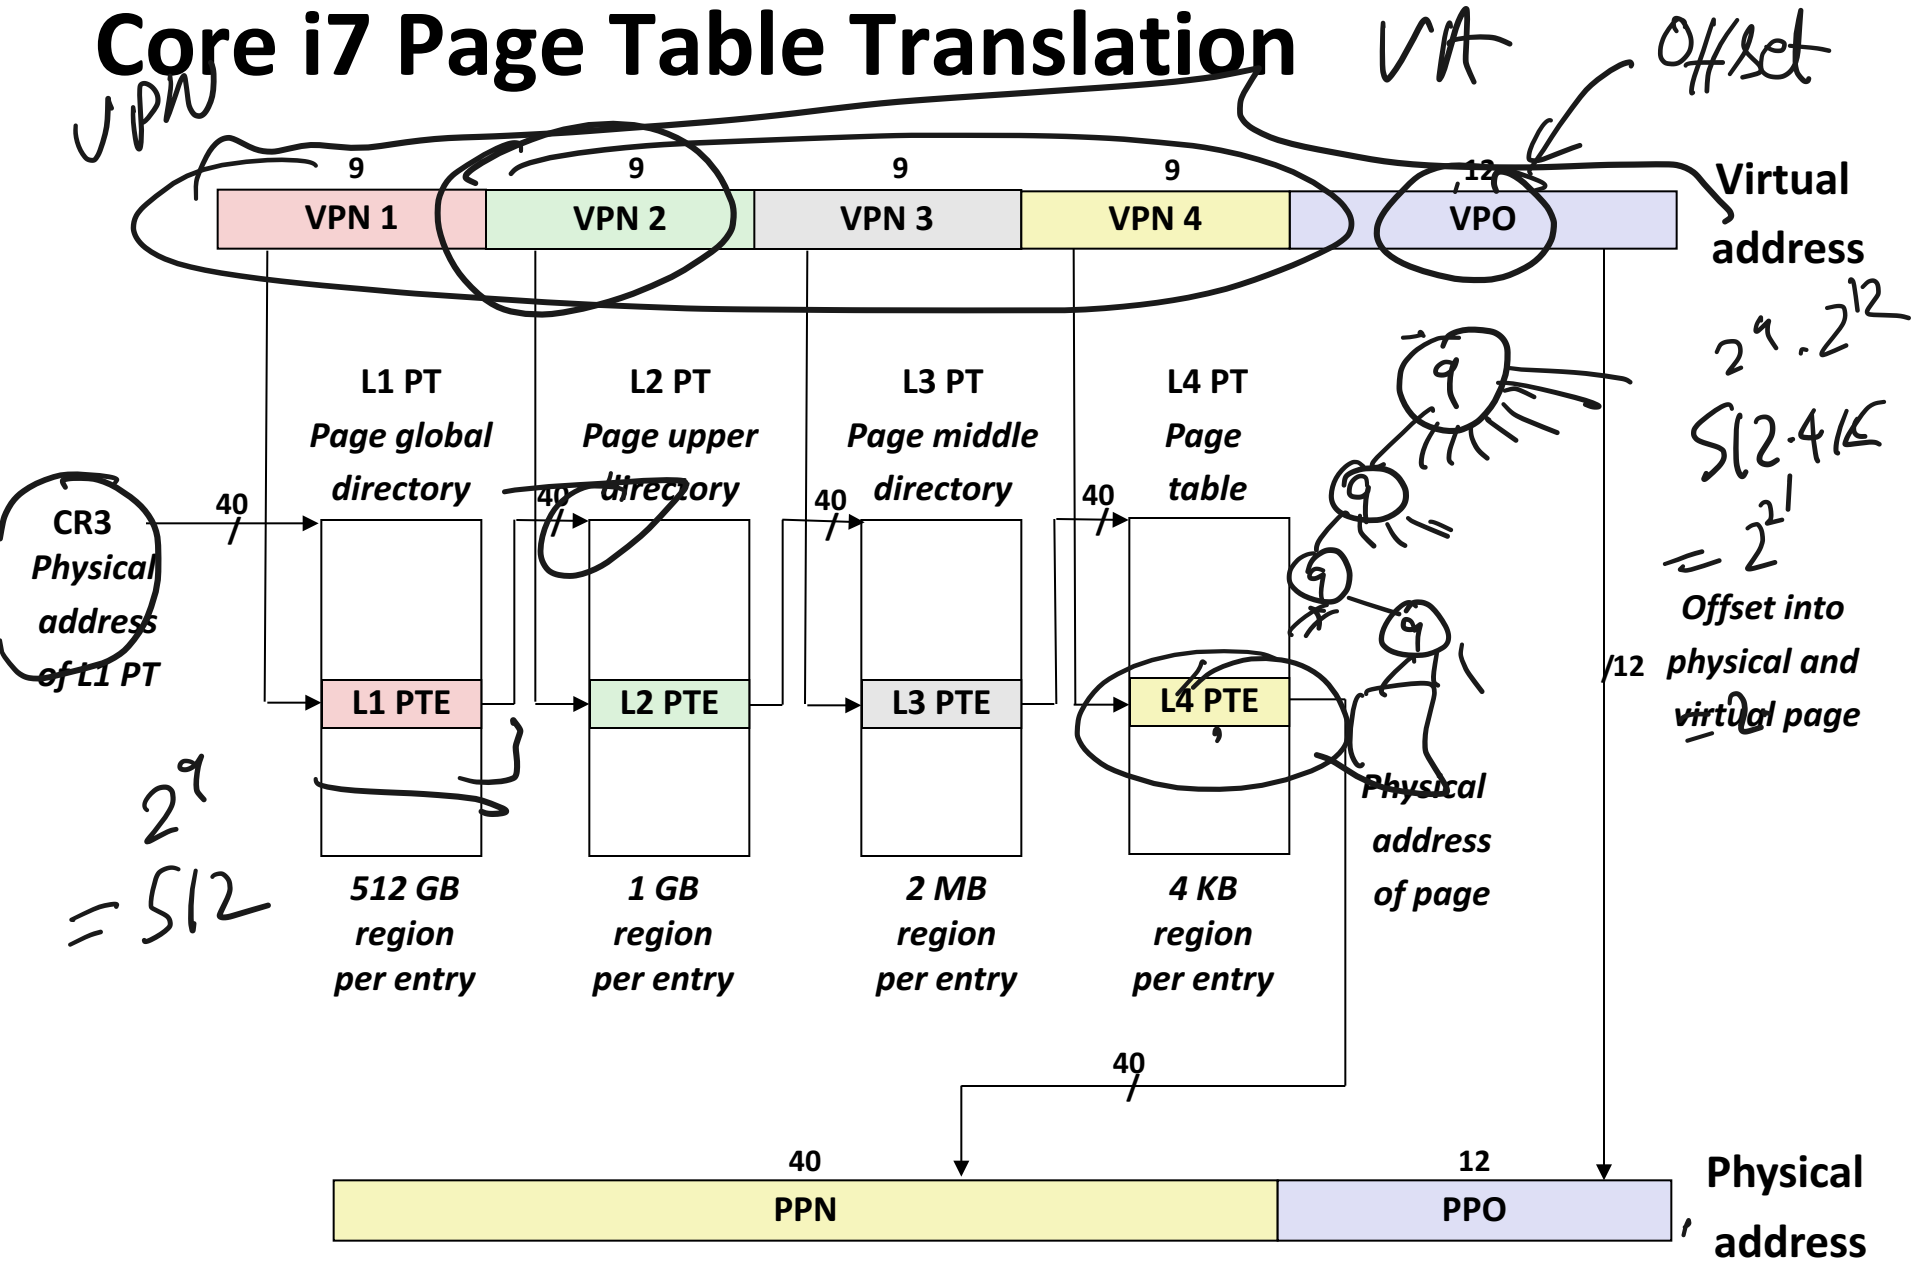
12 bits

$$= 2^{12}$$

$$= 2^2 \cdot 2^{10}$$

$$= 4 \cdot K$$

# Core i7 Page Table Translation





# Core i7 Level 4 Page Table Entries

63	62	52	51	12	11	9	8	7	6	5	4	3	2	1	0
XD	Unused	Page physical base address				Unused	G		D	A	CD	WT	U/S	R/W	P=1
Available for OS (for example, if page location on disk)															P=0

**Each entry references a 4K child page. Significant fields:**

**P:** Child page is present in memory (1) or not (0)

**R/W:** Read-only or read-write access permission for this page

**U/S:** User or supervisor mode access

**WT:** Write-through or write-back cache policy for this page

**A:** Reference bit (set by MMU on reads and writes, cleared by software)


**D:** Dirty bit (set by MMU on writes, cleared by software)

**Page physical base address:** 40 most significant bits of physical page address  
(forces pages to be 4KB aligned)

**XD:** Disable or enable instruction fetches from this page.

# Core i7 Level 1-3 Page Table Entries

VIA 48  
PA 52



Each entry references a 4K child page table. Significant fields:

PPN

**P:** Child page table present in physical memory (1) or not (0).

**R/W:** Read-only or read-write access access permission for all reachable pages.

**U/S:** user or supervisor (kernel) mode access permission for all reachable pages.


**WT:** Write-through or write-back cache policy for the child page table.

**A:** Reference bit (set by MMU on reads and writes, cleared by software).

**PS:** Page size: if bit set, we have 2 MB or 1 GB pages (bit can be set in Level 2 and 3 PTEs only).

**Page table physical base address:** 40 most significant bits of physical page table address (forces page tables to be 4KB aligned)


**XD:** Disable or enable instruction fetches from all pages reachable from this PTE.




- P**      0 The fault was caused by a non-present page.  
           1 The fault was caused by a page-level protection violation.
- W/R**    0 The access causing the fault was a read.  
           1 The access causing the fault was a write.
- U/S**    0 A supervisor-mode access caused the fault.  
           1 A user-mode access caused the fault.
- RSVD**   0 The fault was not caused by reserved bit violation.  
           1 The fault was caused by a reserved bit set to 1 in some  
             paging-structure entry.
- I/D**    0 The fault was not caused by an instruction fetch.  
           1 The fault was caused by an instruction fetch.
- PK**     0 The fault was not caused by protection keys.  
           1 There was a protection-key violation.
- SGX**    0 The fault is not related to SGX.  
           1 The fault resulted from violation of SGX-specific access-control  
             requirements.

**Figure 4-12. Page-Fault Error Code**

# End-to-end Core i7 Address Translation




# Cute Trick for Speeding Up L1 Access



## ■ Observation

- Bits that determine CI identical in virtual and physical address
- Can index into cache while address translation taking place
- Cache carefully sized to make this possible: 64 sets, 64-byte cache blocks
- Means 6 bits for cache index, 6 for *cache* offset
- That's 12 bits; matches *VPO*, *PPO* → One reason pages are  $2^{12}$  bits = 4 KB

# Virtual Address Space of a Linux Process



## Q. PAGES REQUIRED TO ALLOCATE

- 1 BYTE

- 4 KB (PAGE SIZE)  $\equiv 2^{12}$  BYTES

-  $512 \cdot 4 \text{ KB} \equiv 2^9 \cdot 2^{12} \text{ BYTES} \equiv 2 \text{ MB}$

-  $512 \cdot 2 \text{ MB} \equiv 1 \text{ GB} \equiv 2^9 \cdot 2^9 \cdot 2^{12} \text{ BYTES}$