

# 11/14 Endless Paging


## Paging review

- VM  $\Rightarrow$  pages, PM  $\Rightarrow$  frames
- single level vs. multilevel page table
- TLB + HW page table walk
  - $\rightarrow$  architecture spec defines page table format

# X86-64 Page Table Format

page #

12 bits  $2^{12} = 4KB$


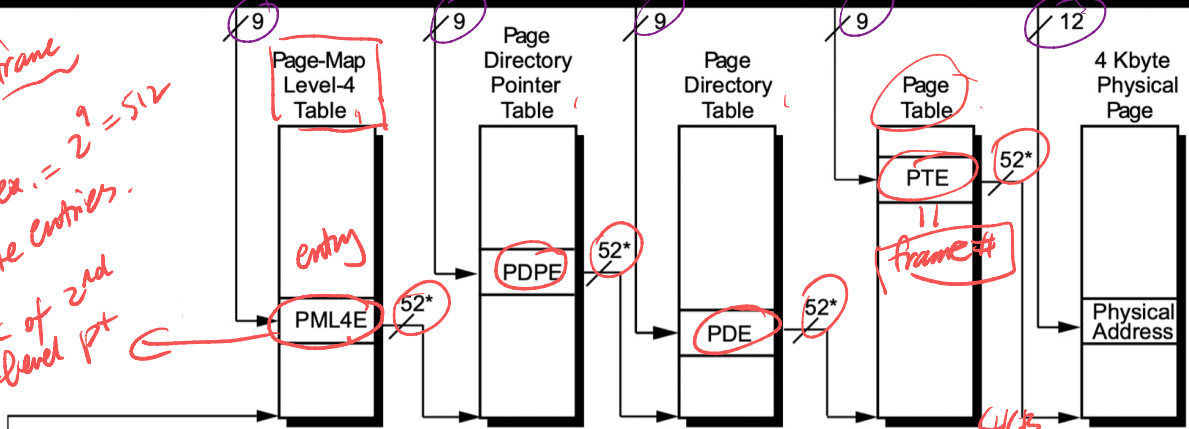


vaddr

frame

4KB PT.  
9 bits to index.  $2^9 = 512$

$\frac{2^{14}}{2^9} = 2^5 = 32$  loc of 2nd level PT



paddr

4KB this is an architectural limit. A given processor implementation may support fewer bits

Let's allocate page table for vaddr 0x1000

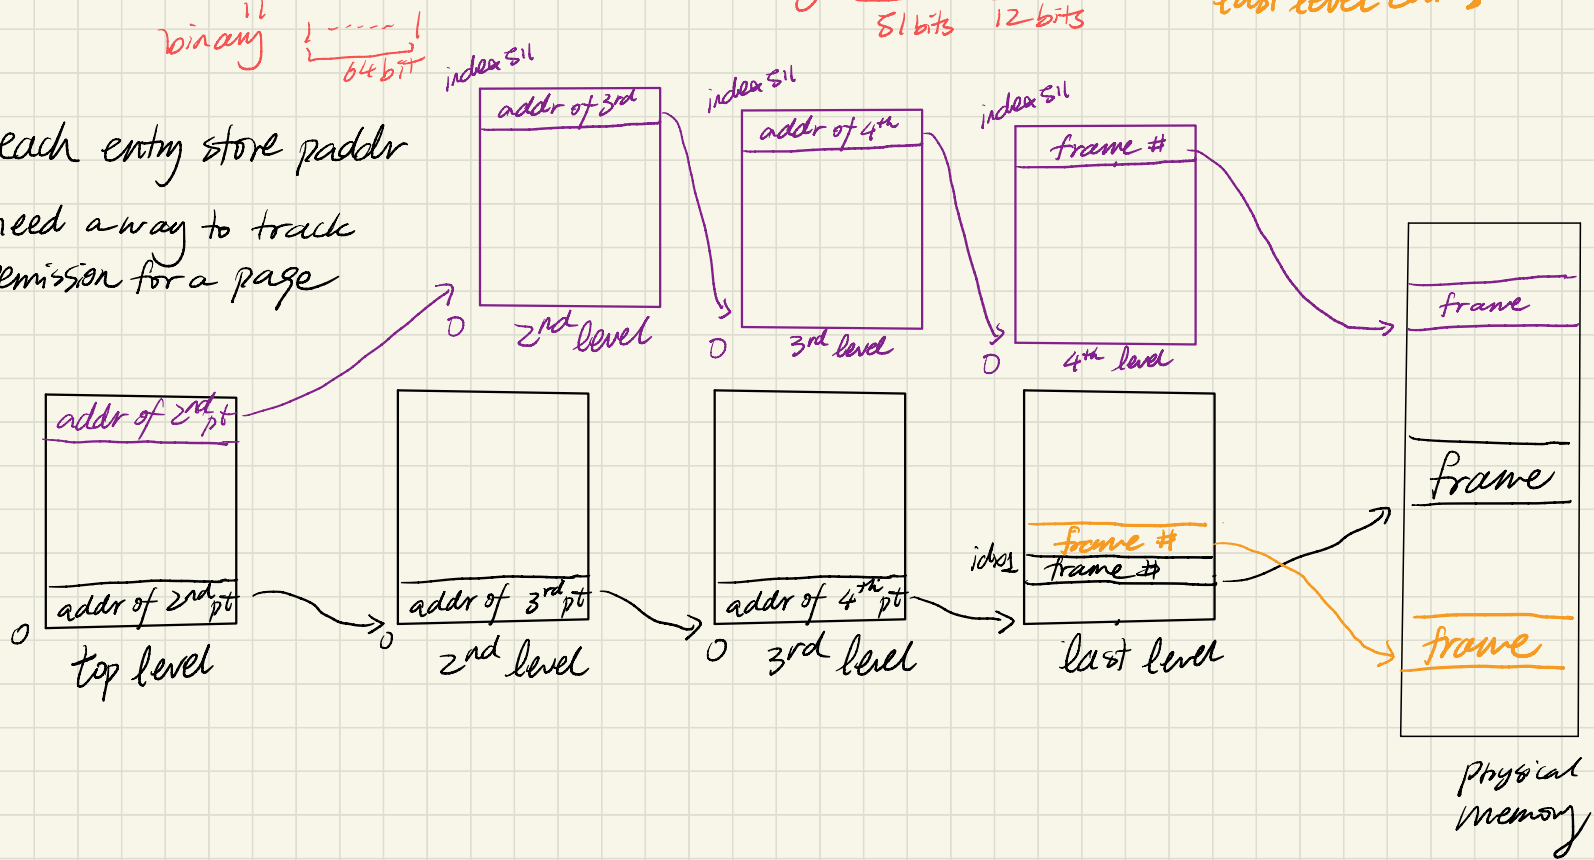
What about 0xfff....

binary <sup>11</sup> 1-----1  
64 bits

binary <sup>11</sup> 0-----0 10-----0  
51 bits 12 bits

What about 0x200?  
→ same PT but different last level entry

- each entry store paddr
- need a way to track permission for a page





What happens when hw encounters error during a page table walk?

- invalid page table entry (present bit == 0)
- permission mismatch (write to a read only page)

★ page fault (exception 14)

→ kernel handles page fault (see trap.c)

→ How?

→ cases

- valid
  - ① stack, heap growth → initialize to 0s (why?)  
(allocate a new frame & create the mapping)
  - ② cow write access  
(allocate a new frame, copy content, update mapping)
  - ③ memory mapped file, swapped pages
- invalid = bad address, actual permission mismatch (user access kernel addr)  
(terminate the process)

Remember TLB?

★ whenever you are changing the translation mapping, you need to flush TLB to get rid of cached, stale permission

xk = vspaceinstall

How does the kernel determine what kind of fault it is?

→ kernel has bookkeeping structures (machine independent) that track information about each page.

\* Machine independent vs. machine dependent page table  
(vspace, vregion, vpage\_info) (x86-64 pt)

info for  
page fault ←

→ info for  
address  
translation

```
42 struct vregion {
43     enum vr_direction dir; // direction of growth
44     uint64_t va_base; // base of the region
45     uint64_t size; // size of region in bytes
46     struct vpi_page *pages; // pointer to array of page_infos
47 };
48
49 struct vspace {
50     struct vregion regions[NREGIONS]; // the regions for a process' virtual space
51     pml4e_t *pgtbl; // process' page table
52 };
```

first level page table loc.

\* vspace invalidate  
→ generates a new  
x86-64 page table from  
vspace data.

```
17 struct vpage_info { information about each page.
18     short used; // whether the page is in use
19     uint64_t ppn; // physical page number
20     short present; // whether the page is in physical memory
21     short writable; // does the page have write permissions
22     // user defined fields
23
24 };
```