Other II	coc of Doging Suppo	<b></b>
Uther U	ses of Paging Suppo	rt
COMP 321		







## Faster, More Efficient Loading of Text on execve()

### Load each page of text <u>on demand</u> from the program file

- The process's page table is built, include PTEs for all pages of text
  - Do not allocate any physical pages for any of the text
  - For each page, valid = 0, protection = read/execute
  - Kernel must remember *why* it set valid = 0
  - Kernel must remember the executable file and where each text page is
  - Return from execve() and start the new program running

### Each individual text page gets loaded on demand in response to page fault

- Allocate a new physical page, PTE valid = 1, read in *only* that one text page
- If text page is chosen as victim, no need to write out the page, and it will be reloaded later on demand the same way when (if) it is next accessed again

COMP 321

5

Copyright © 2025 David B. Johnson

Page 5















# Faster, More Efficient Setting up bss on execve()

#### Allocate and zero out each page of bss on demand

- The process's page table is built, including PTEs for all pages of bss
  - Do not allocate any physical pages for any of the bss
  - For each bss page, PTE valid = 0, protection = read/write
  - Kernel must remember *why* it set valid = 0
  - Return from execve() and start the new program running

### Each individual bss page gets handled on demand in response to page fault

- Allocate a new physical page, PTE valid = 1, zero out only that one bss page
- If such a bss page is chosen as victim, it gets paged out to the *paging file*, the same as the old way, as if it had not been set up on demand
- Technique is called "demand zero" or "zero-fill on demand"

COMP 321

11

Copyright © 2025 David B. Johnson

Page 11



# Faster, More Efficient fork()

### Allocate physical memory for and copy each page only on demand

- The child's page table is built, including PTEs for all pages as in the parent
  - Do not allocate any new physical pages for any of the child's pages
  - The child's PTEs *share* physical pages with the parent (same PFNs)
  - Return from fork() to allow child to run as a new process

### The technique is called "copy-on-write" (COW)

- As long as those shared pages aren't modified, they can be shared
- OK for text pages, but what about data pages, which might get modified?
  - Kernel turns off PTE write protection bit when setting up copy-on-write
  - Attempt to modify page causes hardware exception, and the kernel makes a copy of just that single page and reenables write in PTE protection

COMP 321

Copyright © 2025 David B. Johnson

Page 13

#### 13

### **Copy-on-Write Details**

### For each copy-on-write shared physical page

- The kernel keeps track of the number of times that page is currently shared
  - After a fork, each page is now shared two times (parent and child)
  - But if either process forks again, the count is now 3, and so on
- The kernel disables write access to the page by turning off bit in all PTEs
  - Each sharing process has its own page table and thus its own PTE for it
- If some write is attempted, that will cause a hardware exception
  - If the page is still being shared (count is > 1), the kernel allocates a new physical page, copies the shared page into it, and decrements count
  - The kernel reenables write access in PTE for that page in that process
  - The kernel returns from exception and hardware re-executes instruction

COMP 321

Copyright © 2025 David B. Johnson













Dte mprotect() Kernel Call
Int mprotect(void addr[.len], size\_t len, int prot);
Dtange the protection on a mapping in your virtual address space
Change page protection starting at "addr" for length of "length" invalid
"addr" must be a multiple of PAGESIZE, but "length" need not be
All pages covered in any part by [addr, addr+length) are changed
"prot" should be PROT\_NONE (no access) or the logical "or" of
PROT\_READ – allow read (e.g., LOAD) accesses from the page
PROT\_WRITE – allow write (e.g., STORE) accesses to the page
PROT\_EXEC – allow CPU execution of instructions from the page







