

Operating Systems

Virtual Memory Basics

Peter Lipp, Daniel Gruss

2021-03-04

Table of contents

1. Address Translation

First Idea: Base and Bound

Segmentation

Simple Paging

Multi-level Paging

2. Address Translation on x86 processors

Address Translation

Address Translation

- OS in control of address translation

Address Translation

- OS in control of address translation
- enables number of advanced features

Address Translation

- OS in control of address translation
- enables number of advanced features
- programmers perspective:

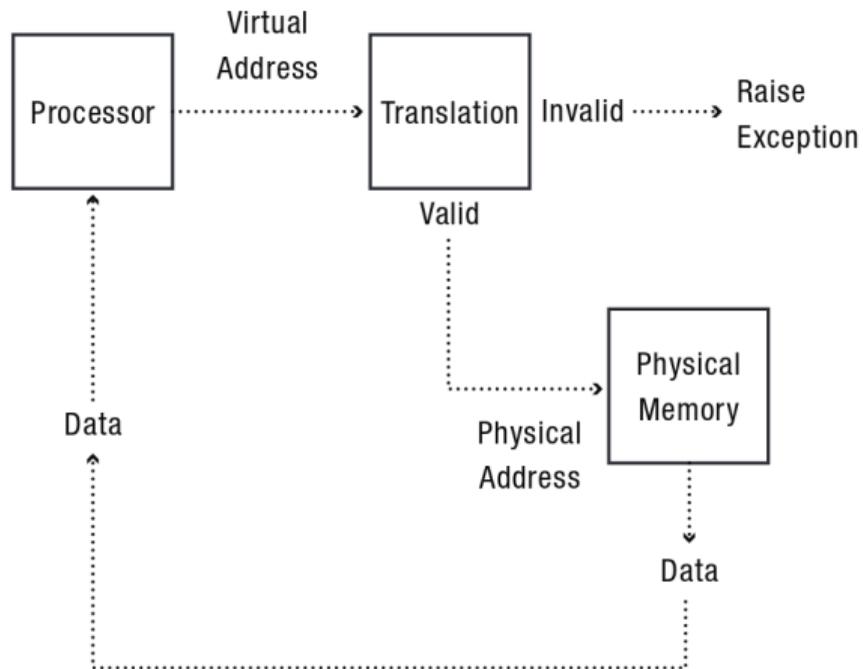
Address Translation

- OS in control of address translation
- enables number of advanced features
- programmers perspective:
 - pointers point to objects etc.

Address Translation

- OS in control of address translation
- enables number of advanced features
- programmers perspective:
 - pointers point to objects etc.
 - transparent: it is not necessary to know how memory reference is converted to data

Address Translation - Idea / Overview



Address Translation Concepts

- Memory protection

Address Translation Concepts

- Memory protection
- Memory sharing

Address Translation Concepts

- Memory protection
- Memory sharing
 - Shared libraries, interprocess communication

Address Translation Concepts

- Memory protection
- Memory sharing
 - Shared libraries, interprocess communication
- Sparse address space

Address Translation Concepts

- Memory protection
- Memory sharing
 - Shared libraries, interprocess communication
- Sparse address space
 - Multiple regions for dynamic allocation (heaps/stacks)

Address Translation Concepts

- Efficiency

Address Translation Concepts

- Efficiency
 - Flexible Memory placement

Address Translation Concepts

- Efficiency
 - Flexible Memory placement
 - Runtime lookup

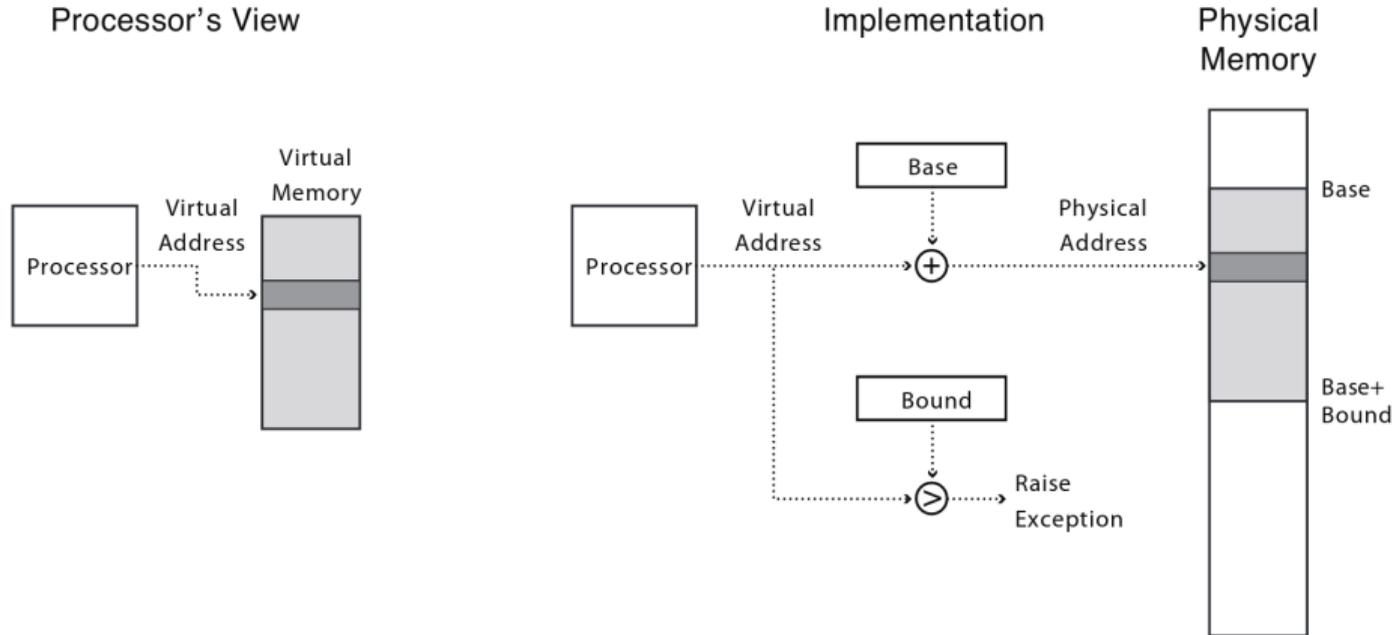
Address Translation Concepts

- Efficiency
 - Flexible Memory placement
 - Runtime lookup
 - Compact translation tables

Address Translation Concepts

- Efficiency
 - Flexible Memory placement
 - Runtime lookup
 - Compact translation tables
- Portability

Base-Limit or Base and bounds



Virtually Addressed Base and Bounds

- Virtual Address: from **0** to an upper **bound**

Virtually Addressed Base and Bounds

- Virtual Address: from **0** to an upper **bound**
- Physical Address: from **base** to **base + bound**

Virtually Addressed Base and Bounds

- Virtual Address: from **0** to an upper **bound**
- Physical Address: from **base** to **base + bound**
- what is saved/restored on a process context switch?

Segmentation

- Small Change: multiple pairs of base-and-bounds registers

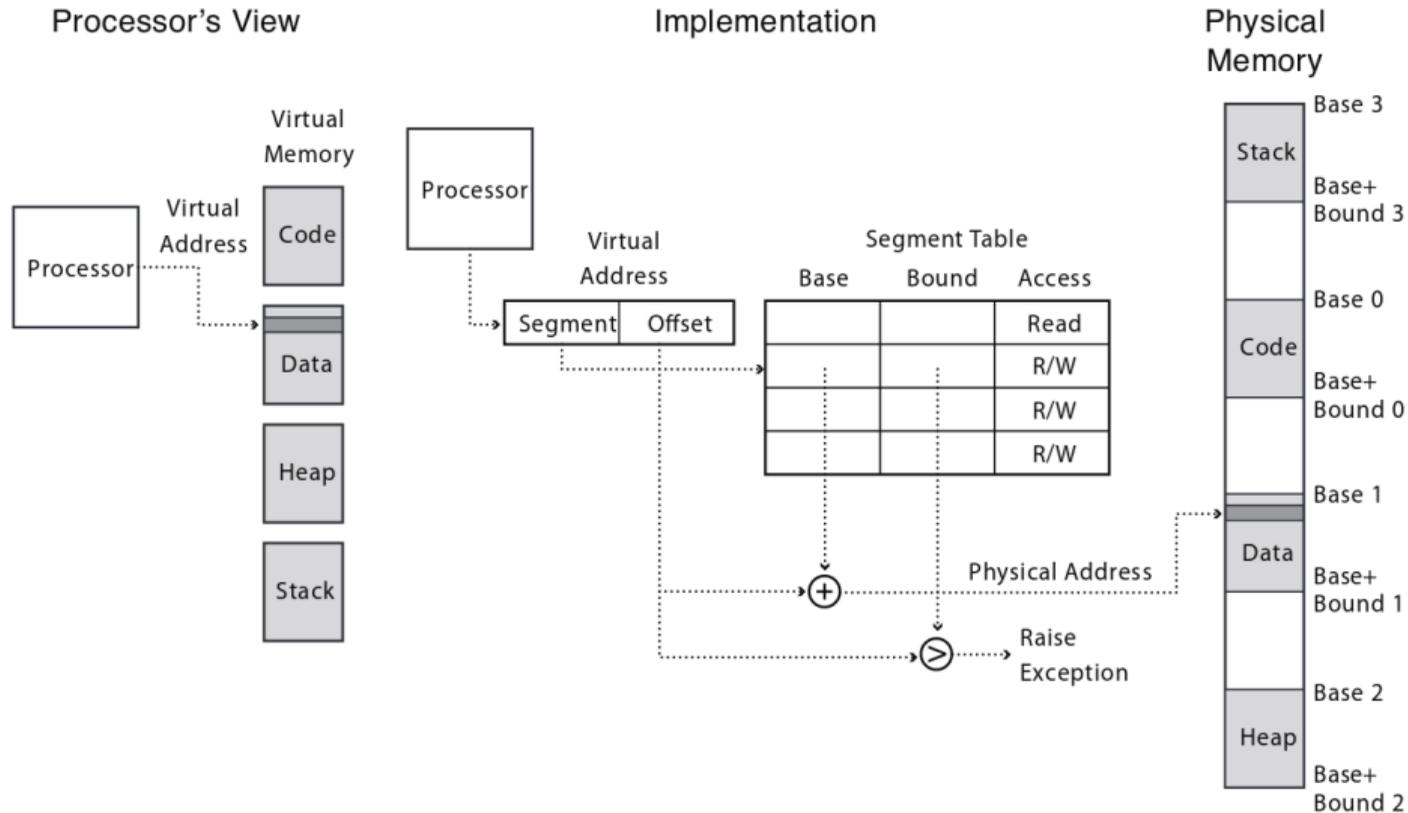
Segmentation

- Small Change: multiple pairs of base-and-bounds registers
- Segmentation

Segmentation

- Small Change: multiple pairs of base-and-bounds registers
- Segmentation
- Each entry controls a portion of the virtual address space

Segmentation



Segmentation

- Segment is a contiguous region of virtual memory

Segmentation

- Segment is a contiguous region of virtual memory
- Each process has a segment table (in hardware)

Segmentation

- Segment is a contiguous region of virtual memory
- Each process has a segment table (in hardware)
 - Entry in table = segment

Segmentation

- Segment is a contiguous region of virtual memory
- Each process has a segment table (in hardware)
 - Entry in table = segment
- Segment can be located anywhere in physical memory

Segmentation

- Segment is a contiguous region of virtual memory
- Each process has a segment table (in hardware)
 - Entry in table = segment
- Segment can be located anywhere in physical memory
 - Each segment has: start, length, access permission

Segmentation

- Segmented Memory has gaps!

Segmentation

- Segmented Memory has gaps!
- no longer contiguous region - set of regions

Segmentation

- Segmented Memory has gaps!
- no longer contiguous region - set of regions
- code and data not adjacent - neither in virtual nor in physical address space

Segmentation

- Segmented Memory has gaps!
- no longer contiguous region - set of regions
- code and data not adjacent - neither in virtual nor in physical address space
- What if: program tries to load data from gap?

Segmentation

- Segmented Memory has gaps!
- no longer contiguous region - set of regions
- code and data not adjacent - neither in virtual nor in physical address space
- What if: program tries to load data from gap?
- Segmentation Fault (trap into OS)

Segmentation

- Segmented Memory has gaps!
- no longer contiguous region - set of regions
- code and data not adjacent - neither in virtual nor in physical address space
- What if: program tries to load data from gap?
- Segmentation Fault (trap into OS)
 - correct programs will not generate references outside valid memory

Segmentation

- Segmented Memory has gaps!
- no longer contiguous region - set of regions
- code and data not adjacent - neither in virtual nor in physical address space
- What if: program tries to load data from gap?
- Segmentation Fault (trap into OS)
 - correct programs will not generate references outside valid memory
 - trying to read or write data that does not exist: bug-indication

Shared Memory

- Processes can share segments

Shared Memory

- Processes can share segments
 - Same start, length, same/different access permissions

Shared Memory

- Processes can share segments
 - Same start, length, same/different access permissions
- Usage:

Shared Memory

- Processes can share segments
 - Same start, length, same/different access permissions
- Usage:
 - sharing code (shared libraries)

Shared Memory

- Processes can share segments
 - Same start, length, same/different access permissions
- Usage:
 - sharing code (shared libraries)
 - interprocess communication

Shared Memory

- Processes can share segments
 - Same start, length, same/different access permissions
- Usage:
 - sharing code (shared libraries)
 - interprocess communication
 - copy on write

Copy on Write

Special kind of shared memory (after fork)

- two processes, both running the same program and almost same data

Copy on Write

Special kind of shared memory (after fork)

- two processes, both running the same program and almost same data
- makes sense not to copy everything

Copy on Write

Special kind of shared memory (after fork)

- two processes, both running the same program and almost same data
- makes sense not to copy everything
- we just need to be made aware if a process writes to a segment and changes the content

Copy on Write

Special kind of shared memory (after fork)

- two processes, both running the same program and almost same data
- makes sense not to copy everything
- we just need to be made aware if a process writes to a segment and changes the content
- reading does not present any problems

Copy on Write

Special kind of shared memory (after fork)

- two processes, both running the same program and almost same data
- makes sense not to copy everything
- we just need to be made aware if a process writes to a segment and changes the content
- reading does not present any problems
- how do we know when a process writes to a segment?

Copy on Write

Special kind of shared memory (after fork)

- two processes, both running the same program and almost same data
- makes sense not to copy everything
- we just need to be made aware if a process writes to a segment and changes the content
- reading does not present any problems
- how do we know when a process writes to a segment?

→ set segment read only

Fork and Copy on Write

Fork:

- Copy segment table into child

Fork and Copy on Write

Fork:

- Copy segment table into child
- Mark parent and child segments read-only

Fork and Copy on Write

Fork:

- Copy segment table into child
- Mark parent and child segments read-only
- Start child process; return to parent

Fork and Copy on Write

Fork:

- Copy segment table into child
- Mark parent and child segments read-only
- Start child process; return to parent

Fork and Copy on Write

Fork:

- Copy segment table into child
- Mark parent and child segments read-only
- Start child process; return to parent

Parent/Child try to write:

- trap into kernel

Fork and Copy on Write

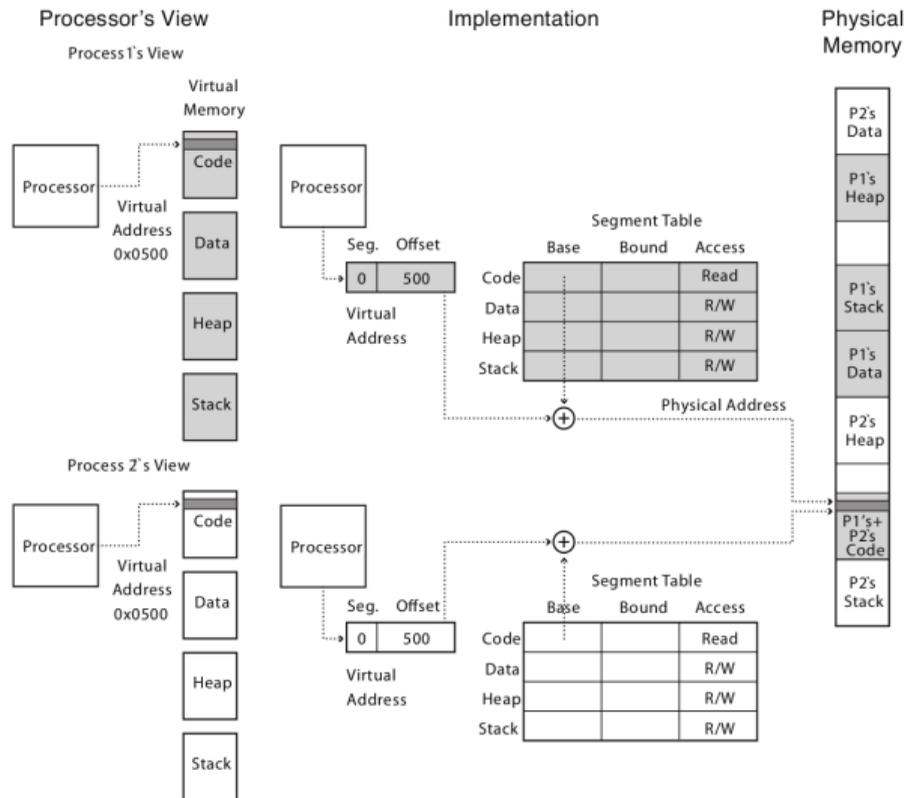
Fork:

- Copy segment table into child
- Mark parent and child segments read-only
- Start child process; return to parent

Parent/Child try to write:

- trap into kernel
- make a copy of the segment and resume

Copy on Write



Zero-on-Reference

- How much physical memory needed for stack or heap?

Zero-on-Reference

- How much physical memory needed for stack or heap?
 - Only what is currently in use

Zero-on-Reference

- How much physical memory needed for stack or heap?
 - Only what is currently in use
- When program uses memory beyond end of stack

Zero-on-Reference

- How much physical memory needed for stack or heap?
 - Only what is currently in use
- When program uses memory beyond end of stack
 - Segmentation fault into OS kernel

Zero-on-Reference

- How much physical memory needed for stack or heap?
 - Only what is currently in use
- When program uses memory beyond end of stack
 - Segmentation fault into OS kernel
 - Kernel allocates some memory

Zero-on-Reference

- How much physical memory needed for stack or heap?
 - Only what is currently in use
- When program uses memory beyond end of stack
 - Segmentation fault into OS kernel
 - Kernel allocates some memory
 - How much?

Zero-on-Reference

- How much physical memory needed for stack or heap?
 - Only what is currently in use
- When program uses memory beyond end of stack
 - Segmentation fault into OS kernel
 - Kernel allocates some memory
 - How much?
 - Zeros the memory

Zero-on-Reference

- How much physical memory needed for stack or heap?
 - Only what is currently in use
- When program uses memory beyond end of stack
 - Segmentation fault into OS kernel
 - Kernel allocates some memory
 - How much?
 - Zeros the memory
 - avoid accidentally leaking information!

Zero-on-Reference

- How much physical memory needed for stack or heap?
 - Only what is currently in use
- When program uses memory beyond end of stack
 - Segmentation fault into OS kernel
 - Kernel allocates some memory
 - How much?
 - Zeros the memory
 - avoid accidentally leaking information!
 - Modify segment table

Zero-on-Reference

- How much physical memory needed for stack or heap?
 - Only what is currently in use
- When program uses memory beyond end of stack
 - Segmentation fault into OS kernel
 - Kernel allocates some memory
 - How much?
 - Zeros the memory
 - avoid accidentally leaking information!
 - Modify segment table
 - Resume process

Paged Translation

- Manage memory in fixed size units, or pages

Paged Translation

- Manage memory in fixed size units, or pages
- Finding a free page is easy

Paged Translation

- Manage memory in fixed size units, or pages
- Finding a free page is easy
 - Bitmap allocation: 0011111100000001100

Paged Translation

- Manage memory in fixed size units, or pages
- Finding a free page is easy
 - Bitmap allocation: 0011111100000001100
 - Each bit represents one physical page number / one physical page frame

Paged Translation

- Manage memory in fixed size units, or pages
- Finding a free page is easy
 - Bitmap allocation: 0011111100000001100
 - Each bit represents one physical page number / one physical page frame
- Each process has its own page table

Paged Translation

- Manage memory in fixed size units, or pages
- Finding a free page is easy
 - Bitmap allocation: 0011111100000001100
 - Each bit represents one physical page number / one physical page frame
- Each process has its own page table
 - Stored in physical memory

Paged Translation

- Manage memory in fixed size units, or pages
- Finding a free page is easy
 - Bitmap allocation: 0011111100000001100
 - Each bit represents one physical page number / one physical page frame
- Each process has its own page table
 - Stored in physical memory
 - Hardware registers

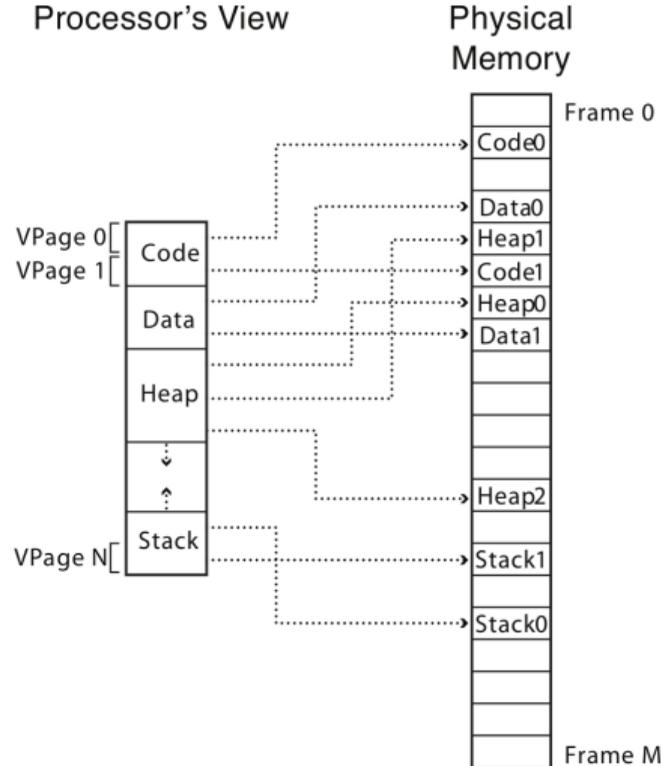
Paged Translation

- Manage memory in fixed size units, or pages
- Finding a free page is easy
 - Bitmap allocation: 0011111100000001100
 - Each bit represents one physical page number / one physical page frame
- Each process has its own page table
 - Stored in physical memory
 - Hardware registers
 - pointer to page table start

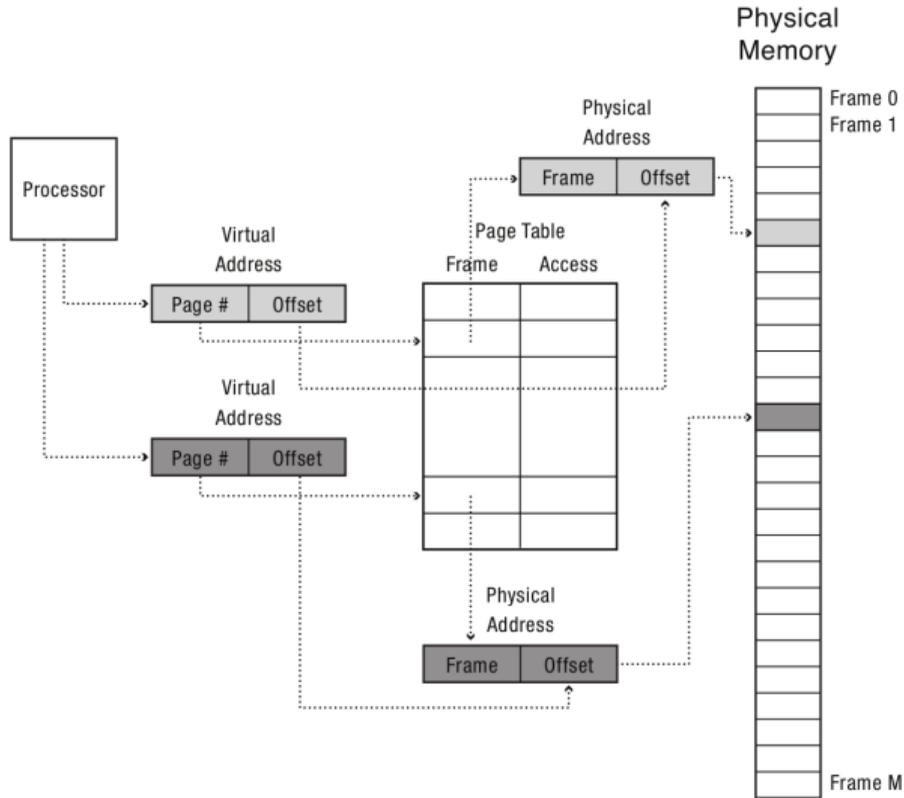
Paged Translation

- Manage memory in fixed size units, or pages
- Finding a free page is easy
 - Bitmap allocation: 0011111100000001100
 - Each bit represents one physical page number / one physical page frame
- Each process has its own page table
 - Stored in physical memory
 - Hardware registers
 - pointer to page table start
 - page table length

Logical View of Page Table Address Translation



paging - implementation



Paging Questions

- With paging, what is saved/restored on a process context switch?

Paging Questions

- With paging, what is saved/restored on a process context switch?
 - Pointer to page table, size of page table

Paging Questions

- With paging, what is saved/restored on a process context switch?
 - Pointer to page table, size of page table
 - Page table itself is in main memory

Paging Questions

- With paging, what is saved/restored on a process context switch?
 - Pointer to page table, size of page table
 - Page table itself is in main memory
- What if page size is very small?

Paging Questions

- With paging, what is saved/restored on a process context switch?
 - Pointer to page table, size of page table
 - Page table itself is in main memory
- What if page size is very small?
- What if page size is very large?

Paging Questions

- With paging, what is saved/restored on a process context switch?
 - Pointer to page table, size of page table
 - Page table itself is in main memory
- What if page size is very small?
- What if page size is very large?
 - Internal fragmentation: if we don't need all of the space inside a fixed size chunk

Paging and Copy on Write

- Can we share pages between processes (similar as segments before)?

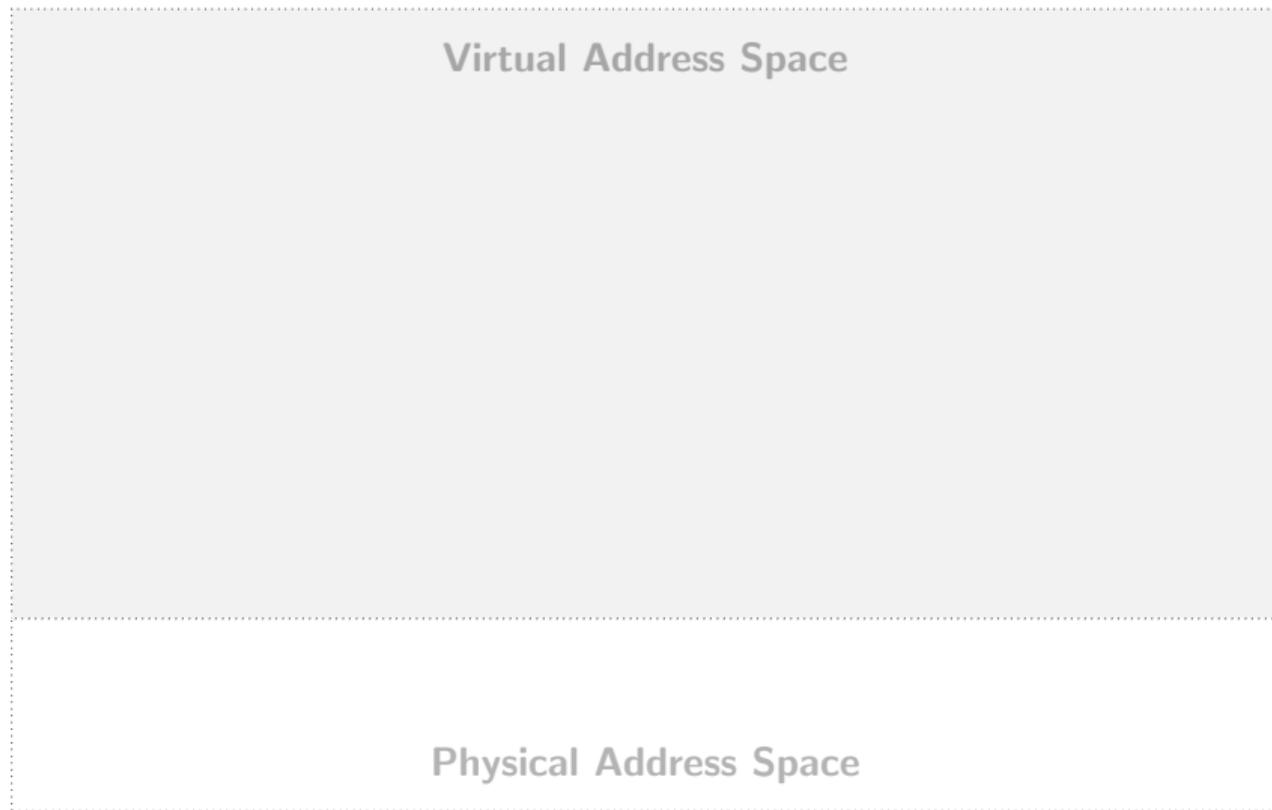
Paging and Copy on Write

- Can we share pages between processes (similar as segments before)?
 - Set entries in both page tables to the same physical page number

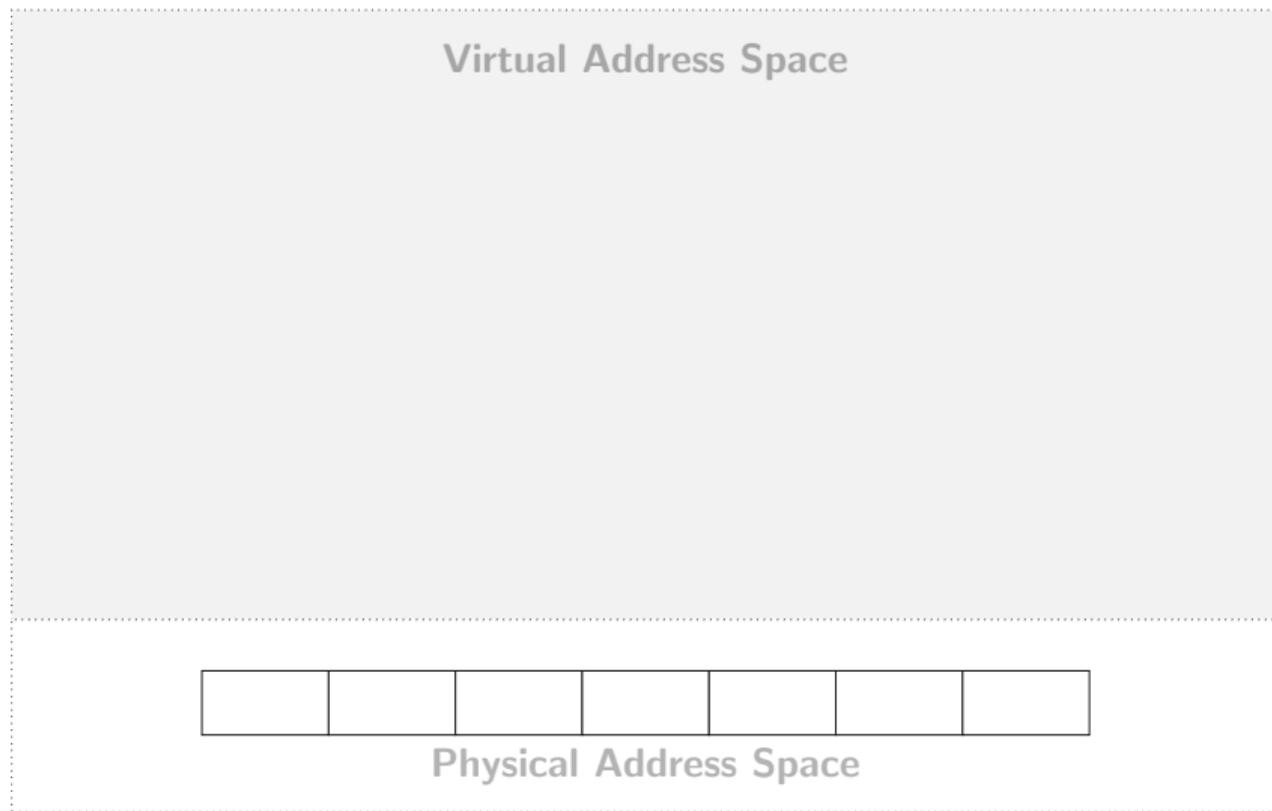
Paging and Copy on Write

- Can we share pages between processes (similar as segments before)?
 - Set entries in both page tables to the same physical page number
 - Need core map of physical page numbers to track which processes are pointing to which physical page numbers (e.g. *reference count*)

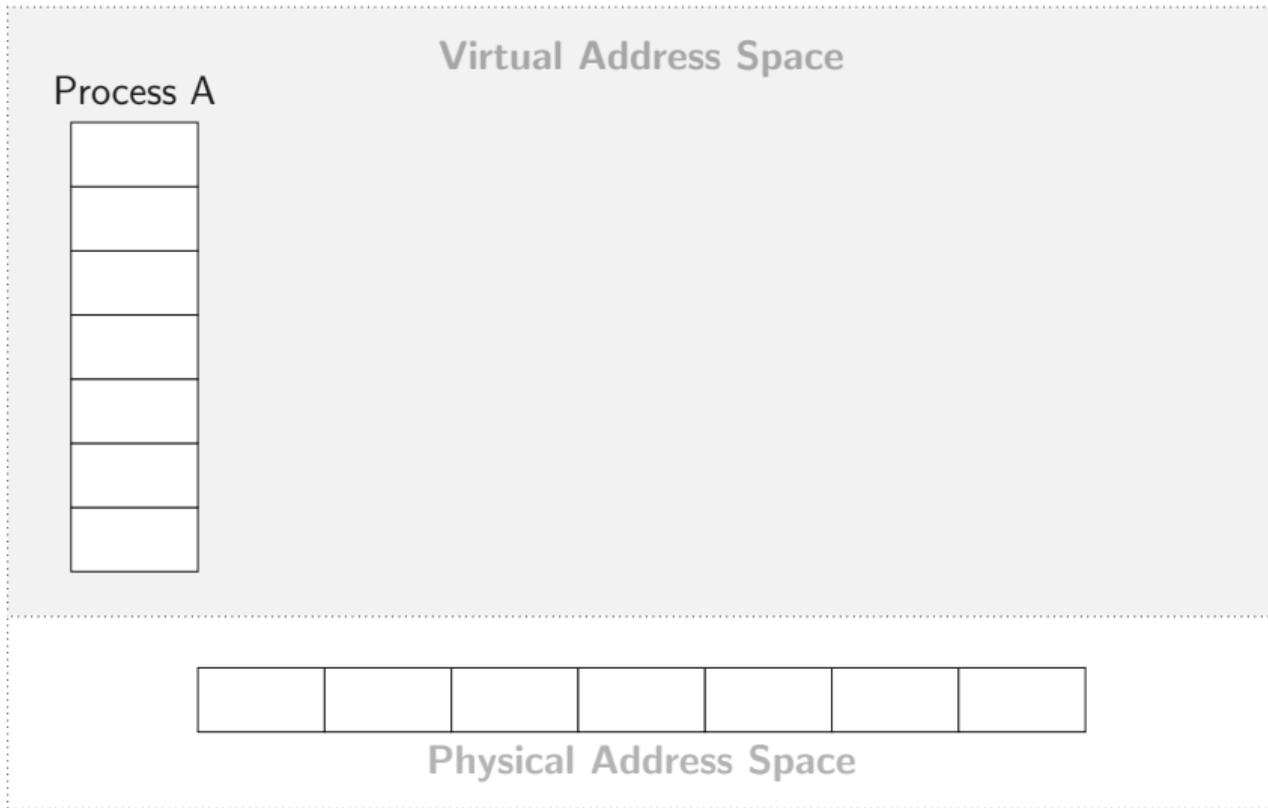
Copy-on-Write on Unix/Linux



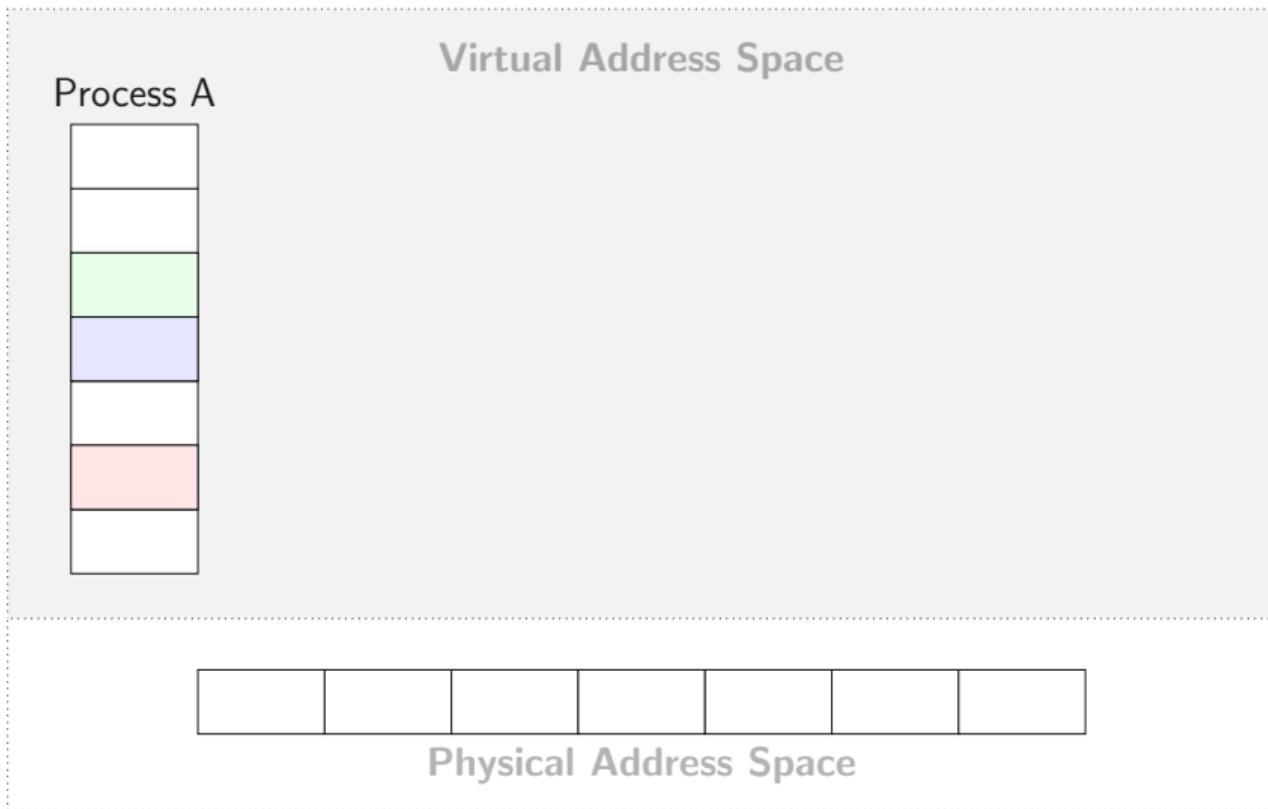
Copy-on-Write on Unix/Linux



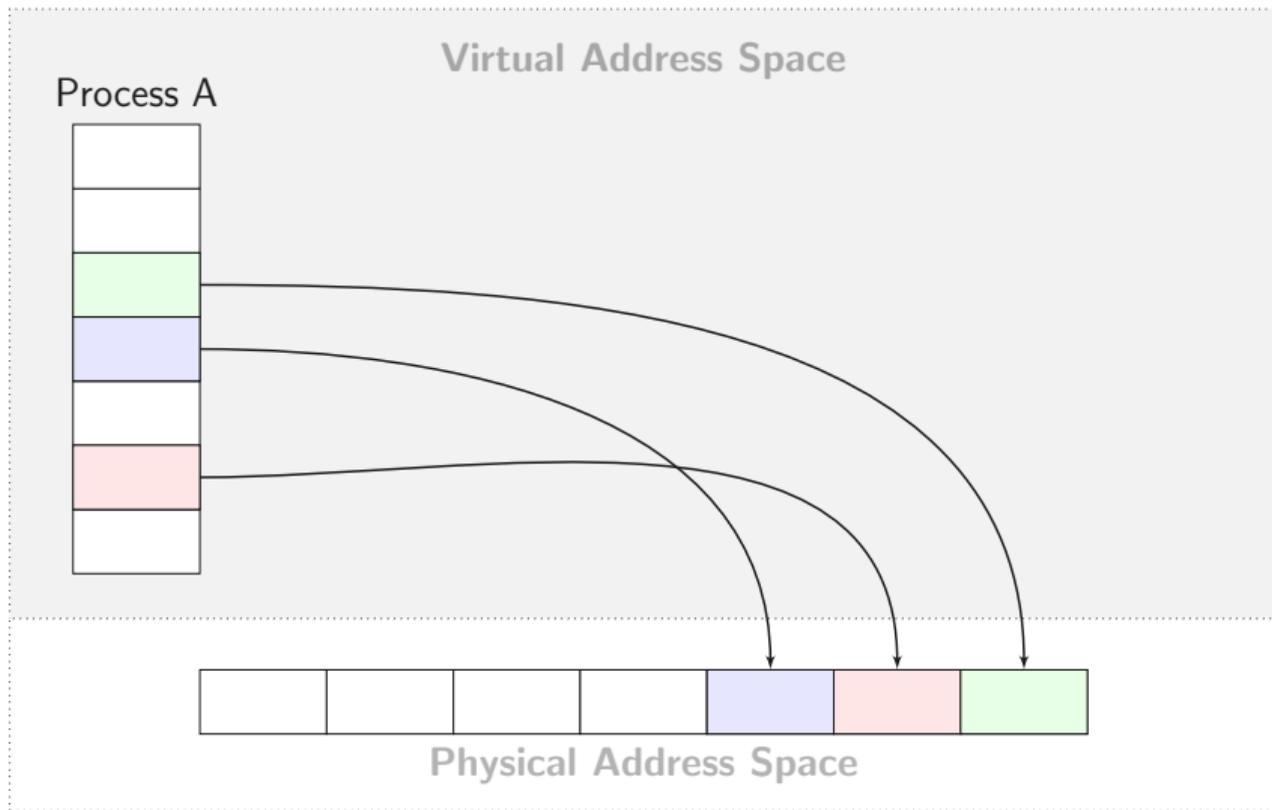
Copy-on-Write on Unix/Linux



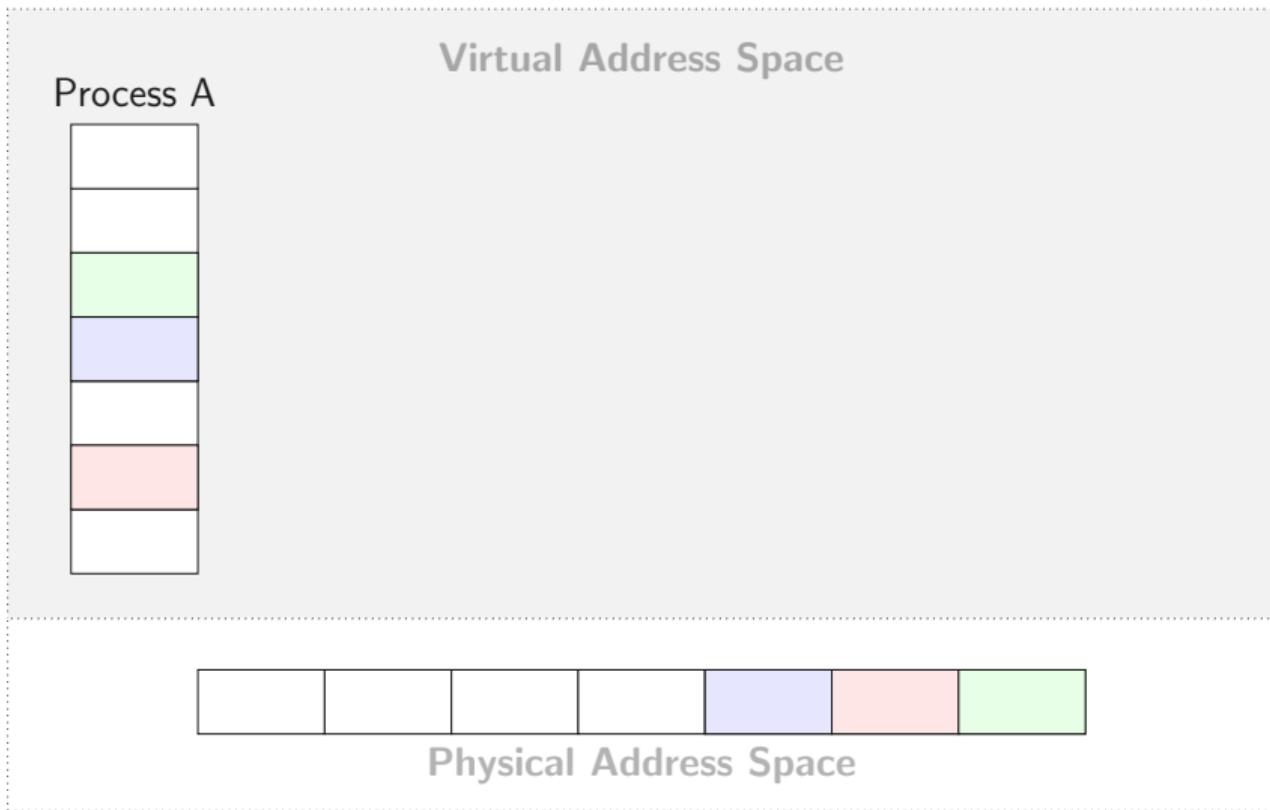
Copy-on-Write on Unix/Linux



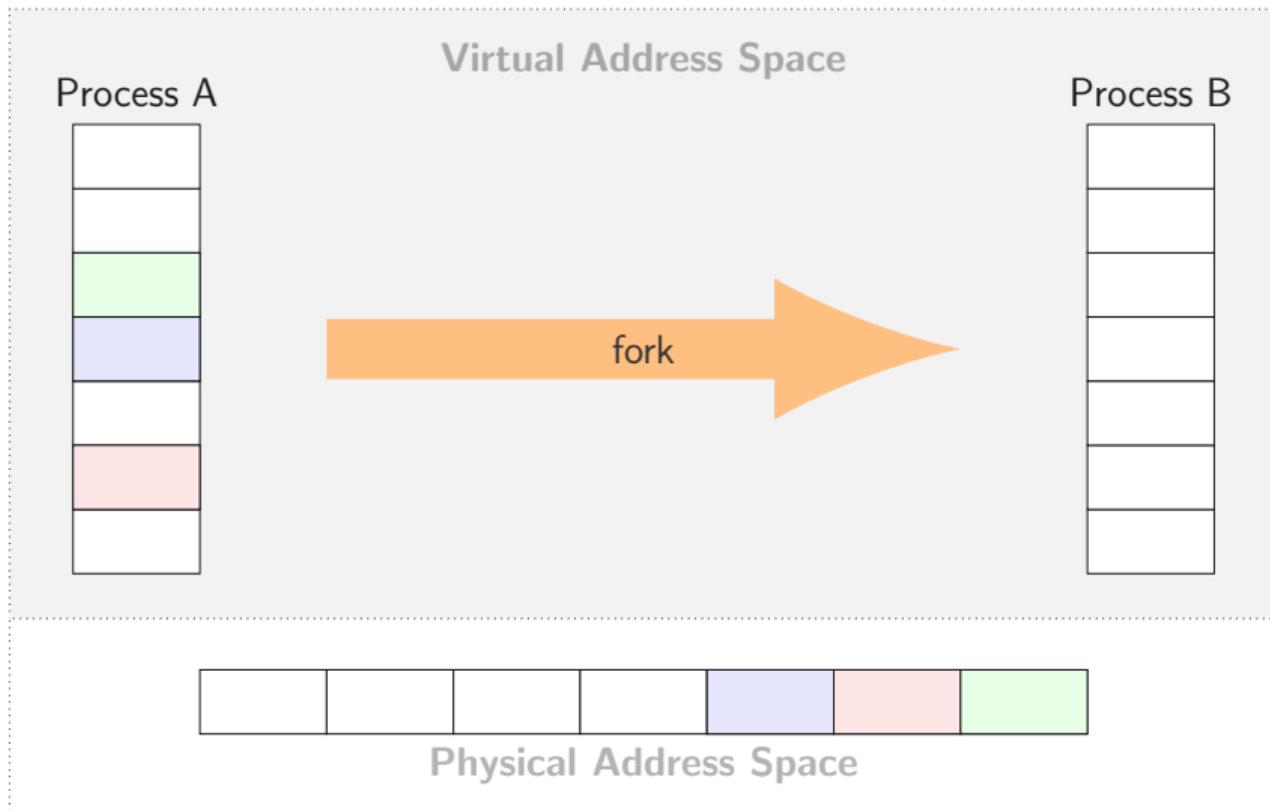
Copy-on-Write on Unix/Linux



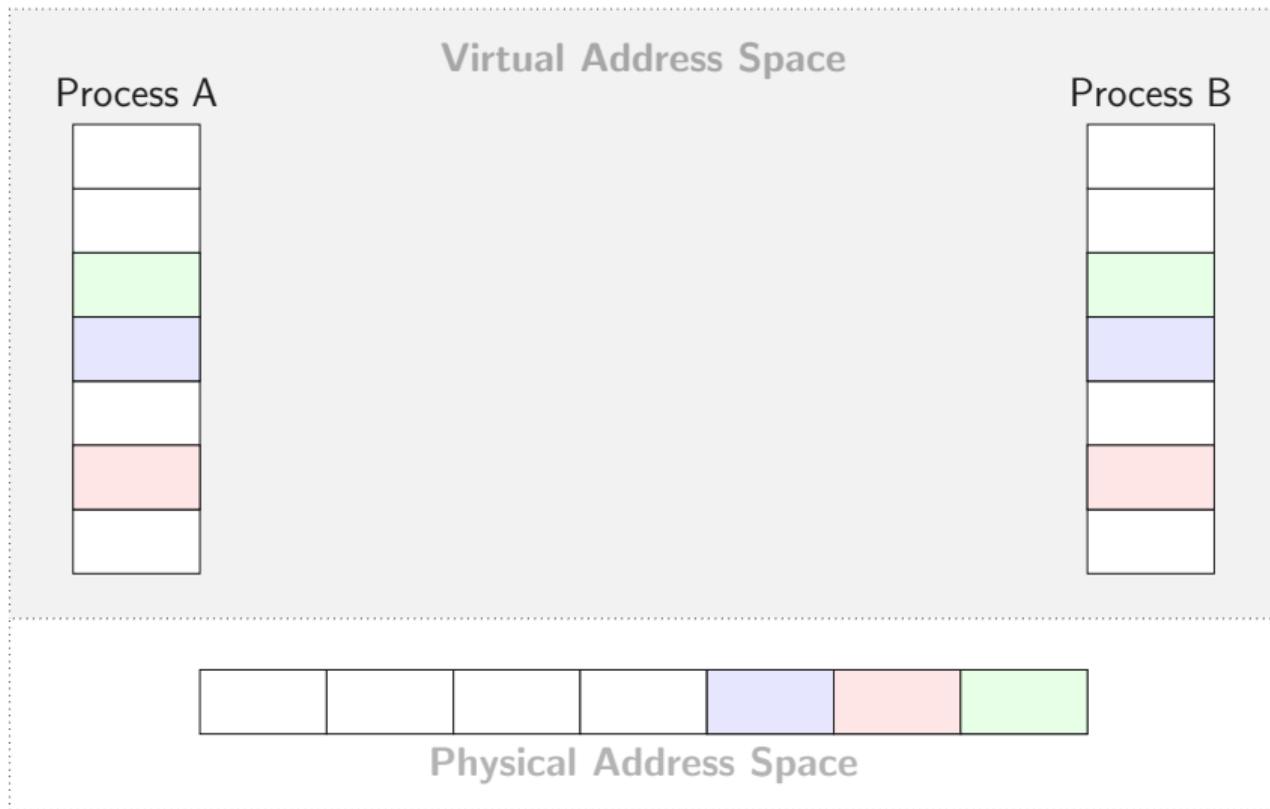
Copy-on-Write on Unix/Linux



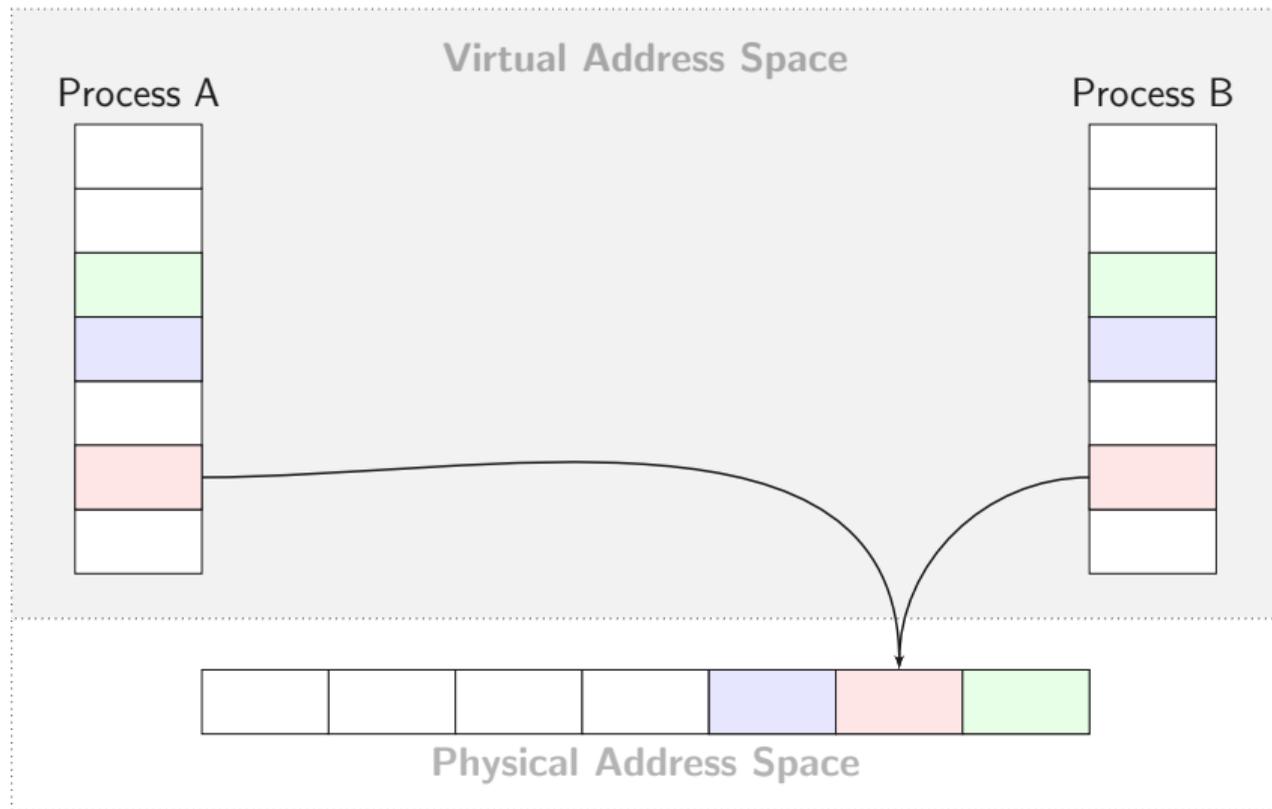
Copy-on-Write on Unix/Linux



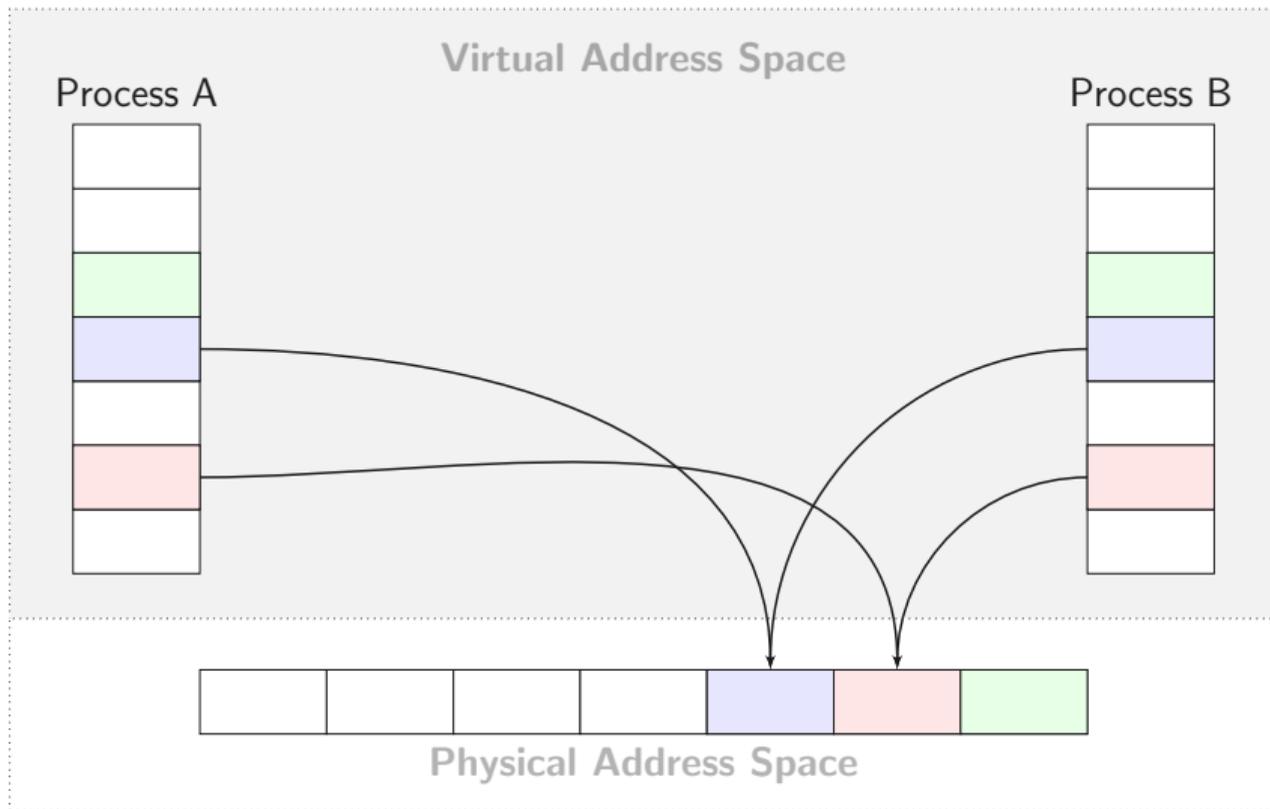
Copy-on-Write on Unix/Linux



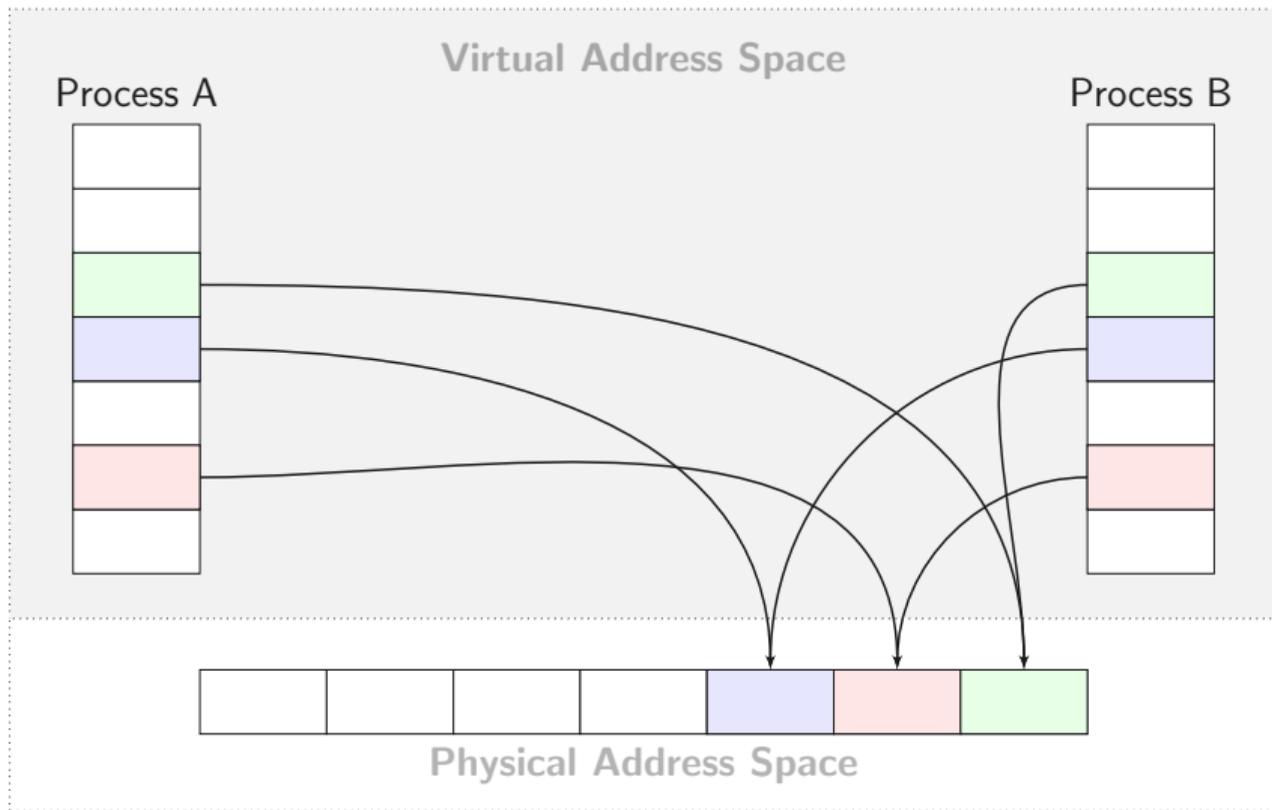
Copy-on-Write on Unix/Linux



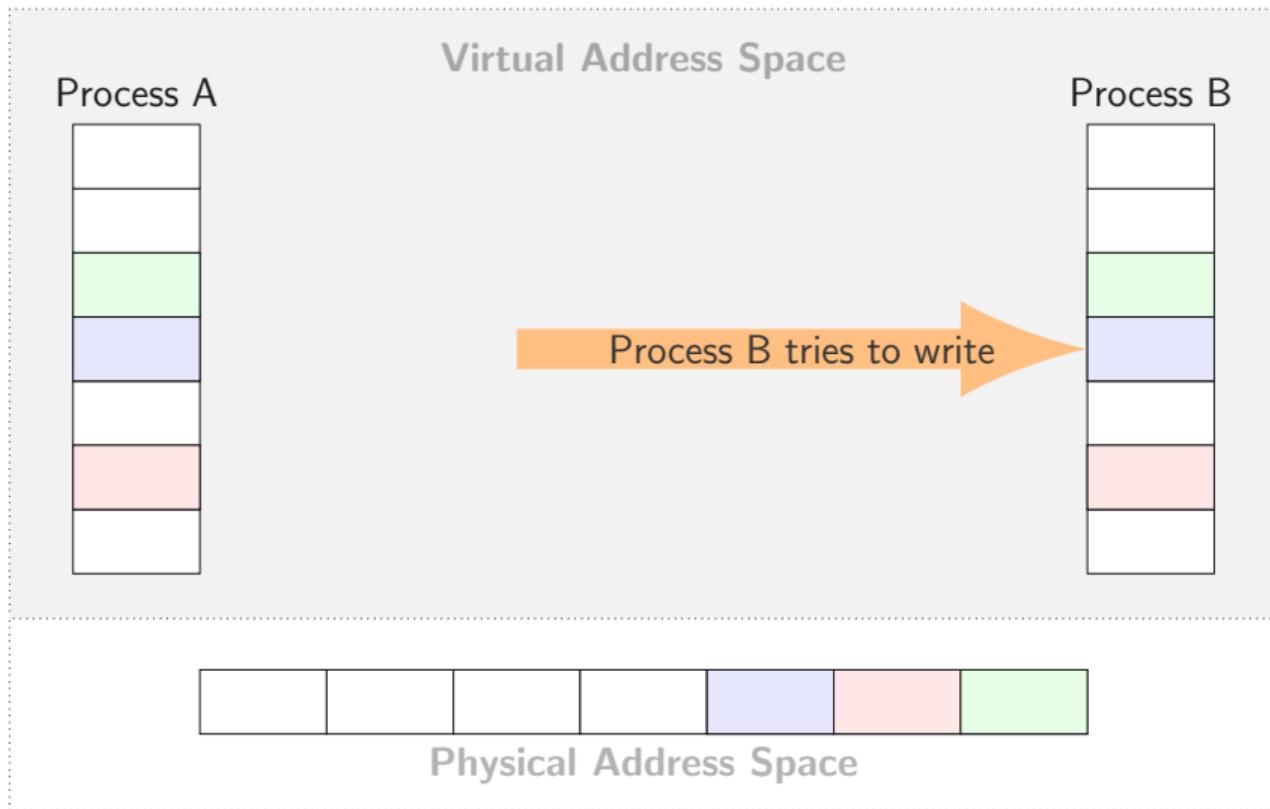
Copy-on-Write on Unix/Linux



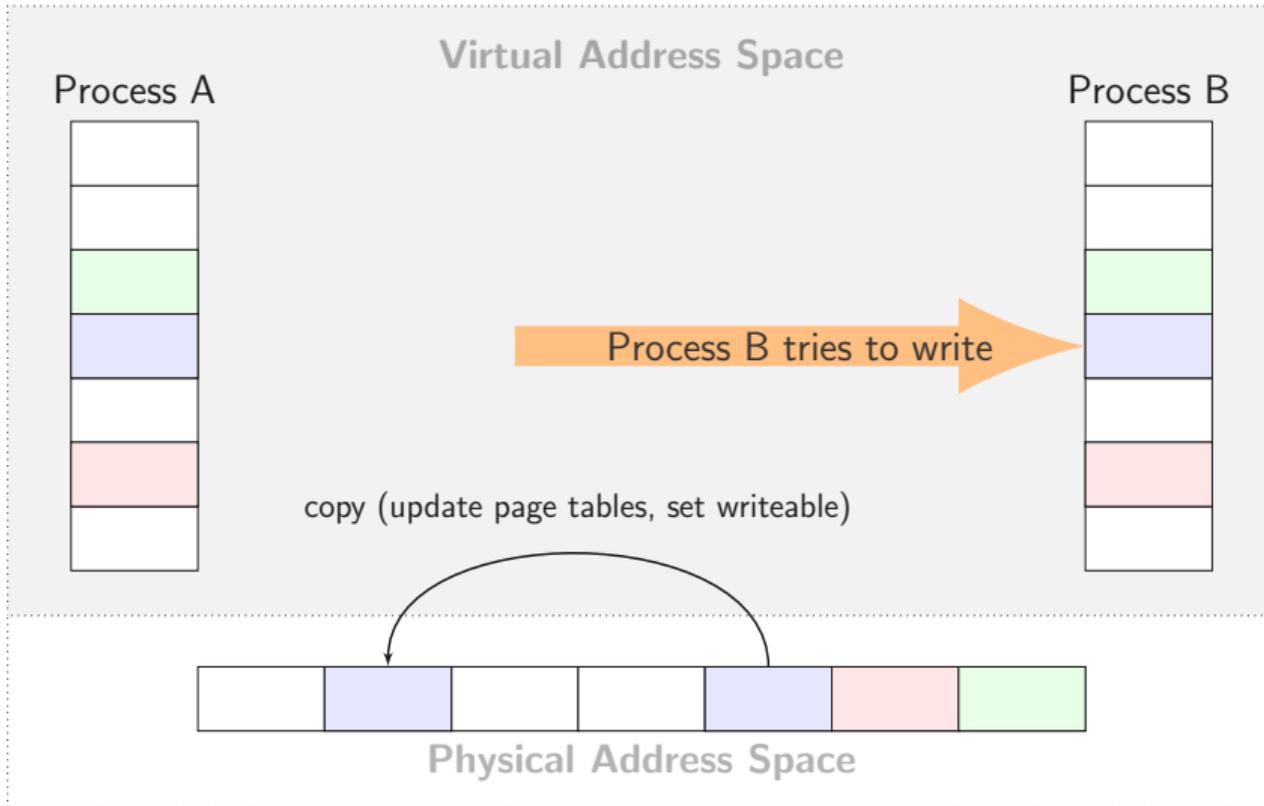
Copy-on-Write on Unix/Linux



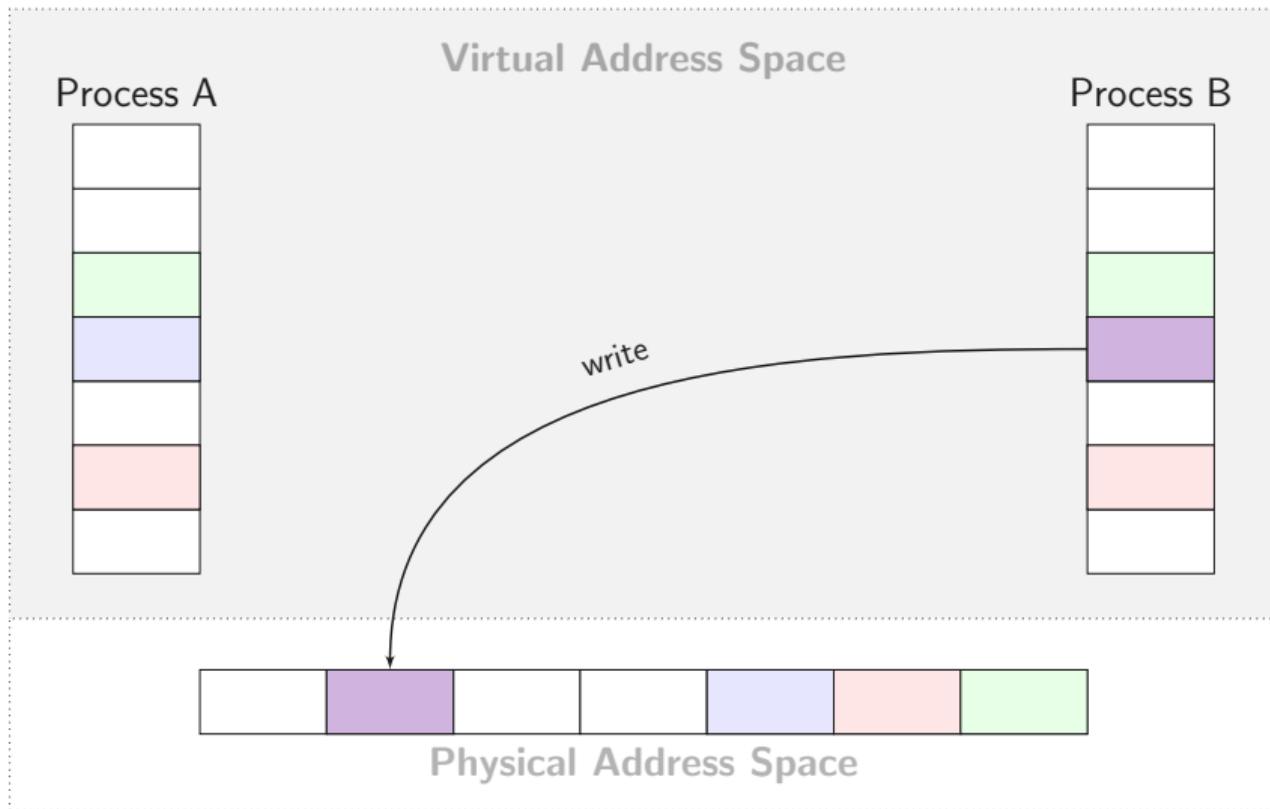
Copy-on-Write on Unix/Linux



Copy-on-Write on Unix/Linux



Copy-on-Write on Unix/Linux



Demand Paging

- Can I start running a program before its code is in physical memory?

Demand Paging

- Can I start running a program before its code is in physical memory?
 - Set all page table entries to invalid

Demand Paging

- Can I start running a program before its code is in physical memory?
 - Set all page table entries to invalid
 - When a page is referenced for first time, kernel trap

Demand Paging

- Can I start running a program before its code is in physical memory?
 - Set all page table entries to invalid
 - When a page is referenced for first time, kernel trap
 - Kernel brings page in from disk

Demand Paging

- Can I start running a program before its code is in physical memory?
 - Set all page table entries to invalid
 - When a page is referenced for first time, kernel trap
 - Kernel brings page in from disk
 - Resume execution

Demand Paging

- Can I start running a program before its code is in physical memory?
 - Set all page table entries to invalid
 - When a page is referenced for first time, kernel trap
 - Kernel brings page in from disk
 - Resume execution
 - Remaining pages can be transferred in the background while program is running

Scheduling a Process (with Demand Paging)

- Only load what's required

Scheduling a Process (with Demand Paging)

- Only load what's required
- Initially start with no pages in memory

Scheduling a Process (with Demand Paging)

- Only load what's required
- Initially start with no pages in memory
- Process will be scheduled eventually. What happens?

Scheduling a Process (with Demand Paging)

- Only load what's required
- Initially start with no pages in memory
- Process will be scheduled eventually. What happens?
 - a page fault will occur when fetching the first instruction

Scheduling a Process (with Demand Paging)

- Only load what's required
- Initially start with no pages in memory
- Process will be scheduled eventually. What happens?
 - a page fault will occur when fetching the first instruction
 - further page faults for stacks and data

Scheduling a Process (with Demand Paging)

- Only load what's required
- Initially start with no pages in memory
- Process will be scheduled eventually. What happens?
 - a page fault will occur when fetching the first instruction
 - further page faults for stacks and data
 - after a while, things will stabilize

Scheduling a Process (with Demand Paging)

- Only load what's required
- Initially start with no pages in memory
- Process will be scheduled eventually. What happens?
 - a page fault will occur when fetching the first instruction
 - further page faults for stacks and data
 - after a while, things will stabilize
- The principle of locality ensures that

Prepaging

Prepaging as an optimization

- If it is known upon scheduling which pages will be required ...

Prepaging

Prepaging as an optimization

- If it is known upon scheduling which pages will be required ...
 - page referenced by instruction pointer, stack pointer, etc.

Prepaging

Prepaging as an optimization

- If it is known upon scheduling which pages will be required ...
 - page referenced by instruction pointer, stack pointer, etc.
- ... load required pages into RAM ahead of time

Prepaging

Prepaging as an optimization

- If it is known upon scheduling which pages will be required ...
 - page referenced by instruction pointer, stack pointer, etc.
- ... load required pages into RAM ahead of time

→ may lower page fault frequency

Sparse Address Spaces

- Every process needs an address space.

Sparse Address Spaces

- Every process needs an address space.
- What if virtual address space is large?

Sparse Address Spaces

- Every process needs an address space.
- What if virtual address space is large?
 - 32-bits, 4KB pages → 1 million page table entries

Sparse Address Spaces

- Every process needs an address space.
- What if virtual address space is large?
 - 32-bits, 4KB pages → 1 million page table entries
 - 64-bits → 4 quadrillion page table entries

Multi-level Translation

- Tree of translation tables

Multi-level Translation

- Tree of translation tables
 - Paged segmentation

Multi-level Translation

- Tree of translation tables
 - Paged segmentation
 - Multi-level page tables

Multi-level Translation

- Tree of translation tables
 - Paged segmentation
 - Multi-level page tables
 - Multi-level paged segmentation

Multi-level Translation

- Fixed-size page as lowest level unit of allocation

Multi-level Translation

- Fixed-size page as lowest level unit of allocation
 - Efficient memory allocation (compared to segments)

Multi-level Translation

- Fixed-size page as lowest level unit of allocation
 - Efficient memory allocation (compared to segments)
 - Efficient for sparse translation tree (compared to simple paging)

Multi-level Translation

- Fixed-size page as lowest level unit of allocation
 - Efficient memory allocation (compared to segments)
 - Efficient for sparse translation tree (compared to simple paging)
 - Efficient disk transfers (fixed size units, page size multiple of disk sector)

Multi-level Translation

- Fixed-size page as lowest level unit of allocation
 - Efficient memory allocation (compared to segments)
 - Efficient for sparse translation tree (compared to simple paging)
 - Efficient disk transfers (fixed size units, page size multiple of disk sector)
 - Easier to build translation lookaside buffers

Multi-level Translation

- Fixed-size page as lowest level unit of allocation
 - Efficient memory allocation (compared to segments)
 - Efficient for sparse translation tree (compared to simple paging)
 - Efficient disk transfers (fixed size units, page size multiple of disk sector)
 - Easier to build translation lookaside buffers
 - Efficient reverse lookup (from physical \rightarrow virtual)

Multi-level Translation

- Fixed-size page as lowest level unit of allocation
 - Efficient memory allocation (compared to segments)
 - Efficient for sparse translation tree (compared to simple paging)
 - Efficient disk transfers (fixed size units, page size multiple of disk sector)
 - Easier to build translation lookaside buffers
 - Efficient reverse lookup (from physical \rightarrow virtual)
 - Fine granularity for protection/sharing

Paged Segmentation

- Process memory is segmented

Paged Segmentation

- Process memory is segmented
- Segment table entry:

Paged Segmentation

- Process memory is segmented
- Segment table entry:
 - Pointer to page table

Paged Segmentation

- Process memory is segmented
- Segment table entry:
 - Pointer to page table
 - Page table length (# of pages in segment)

Paged Segmentation

- Process memory is segmented
- Segment table entry:
 - Pointer to page table
 - Page table length (# of pages in segment)
 - Access permissions

Paged Segmentation

- Process memory is segmented
- Segment table entry:
 - Pointer to page table
 - Page table length (# of pages in segment)
 - Access permissions
- Page table entry:

Paged Segmentation

- Process memory is segmented
- Segment table entry:
 - Pointer to page table
 - Page table length (# of pages in segment)
 - Access permissions
- Page table entry:
 - Physical page number

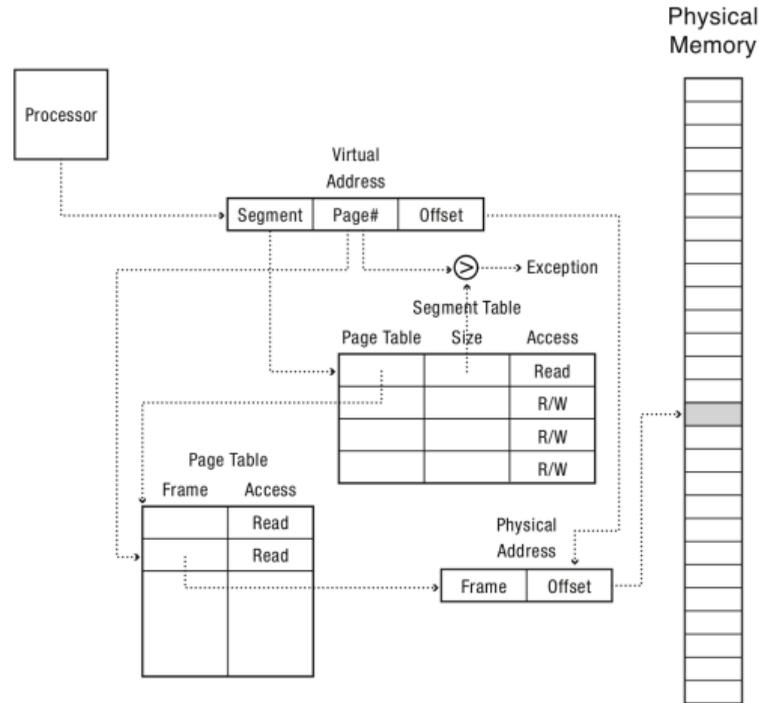
Paged Segmentation

- Process memory is segmented
- Segment table entry:
 - Pointer to page table
 - Page table length (# of pages in segment)
 - Access permissions
- Page table entry:
 - Physical page number
 - Access permissions

Paged Segmentation

- Process memory is segmented
- Segment table entry:
 - Pointer to page table
 - Page table length (# of pages in segment)
 - Access permissions
- Page table entry:
 - Physical page number
 - Access permissions
- Share/protection at either page or segment-level

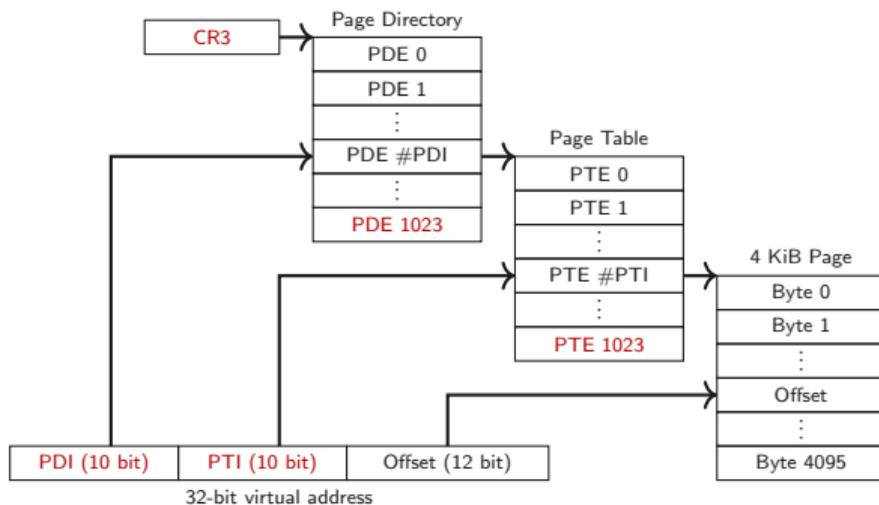
Paged Segmentation



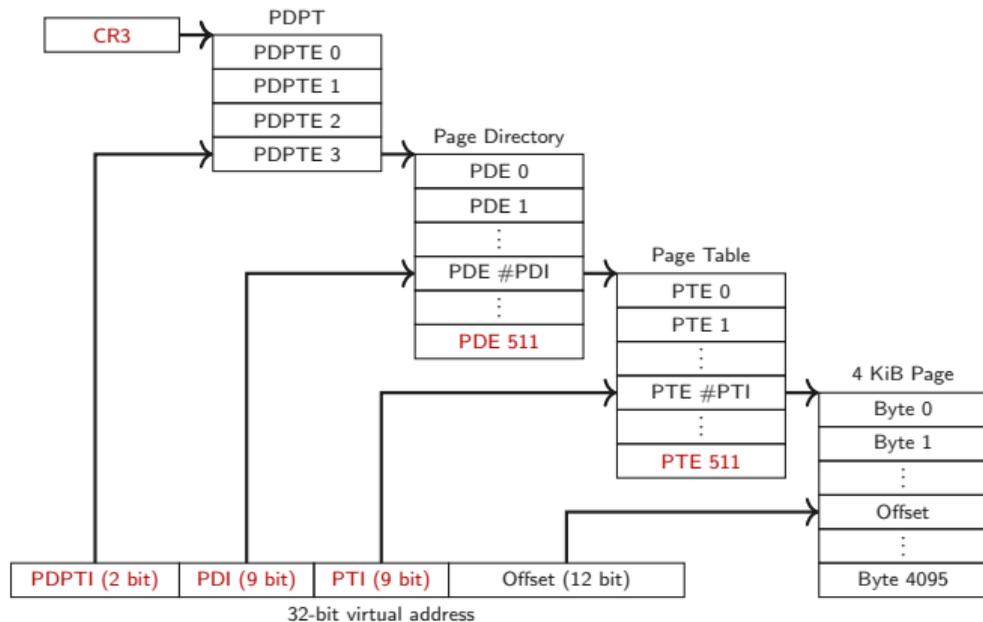
Question

- With paged segmentation, what must be saved/restored across a process context switch?

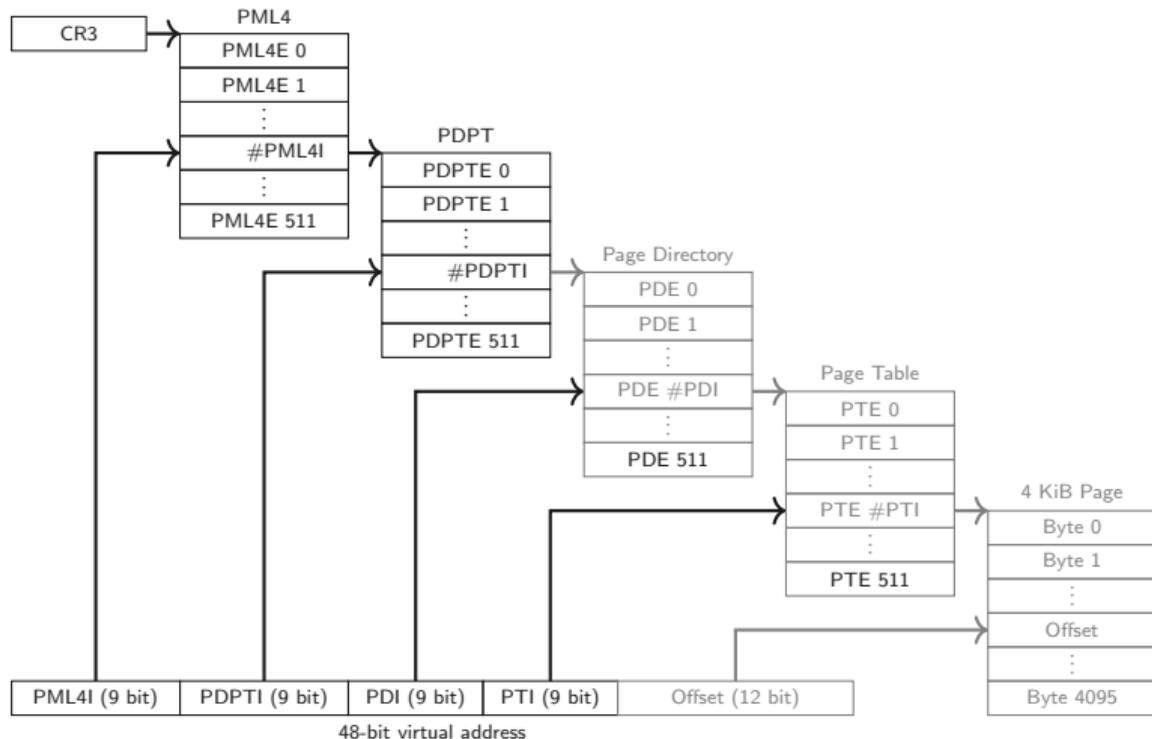
Paging: x86-32 with page size 4 KiB



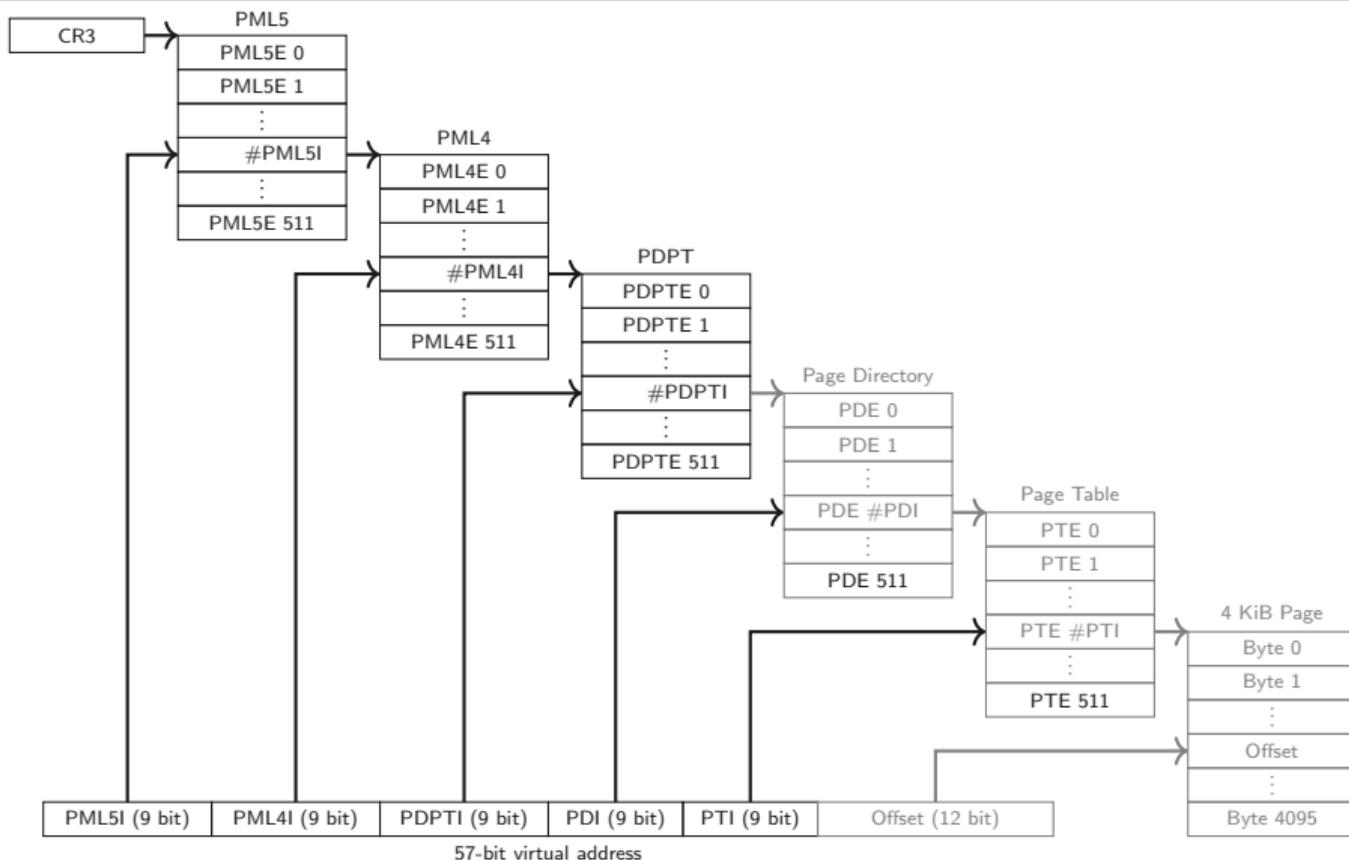
Paging: x86-32-PAE with page size 4 KiB



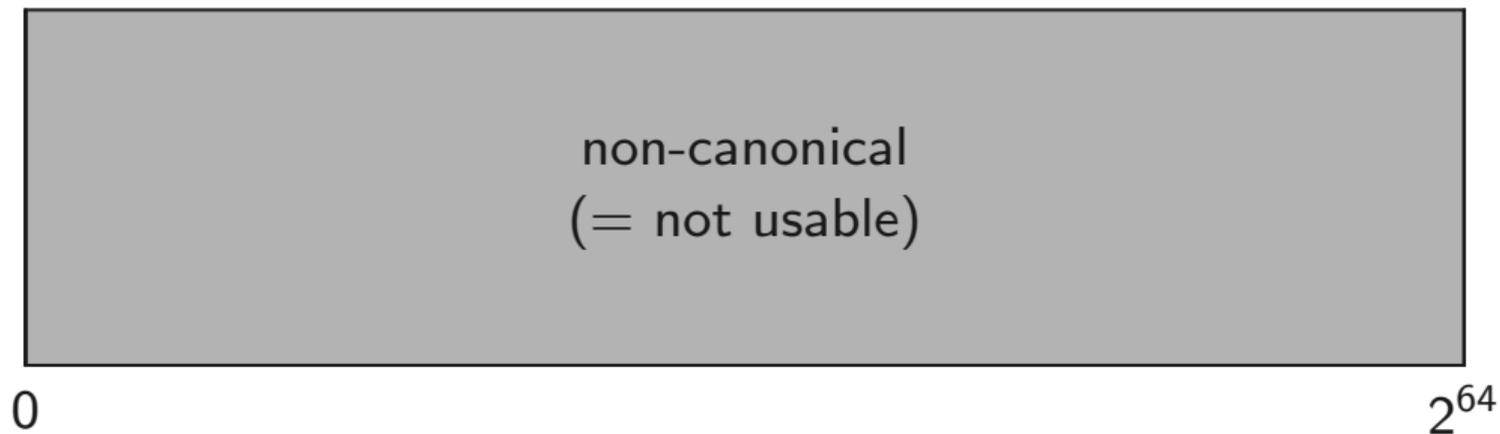
Paging: x86-64 with page size 4 KiB



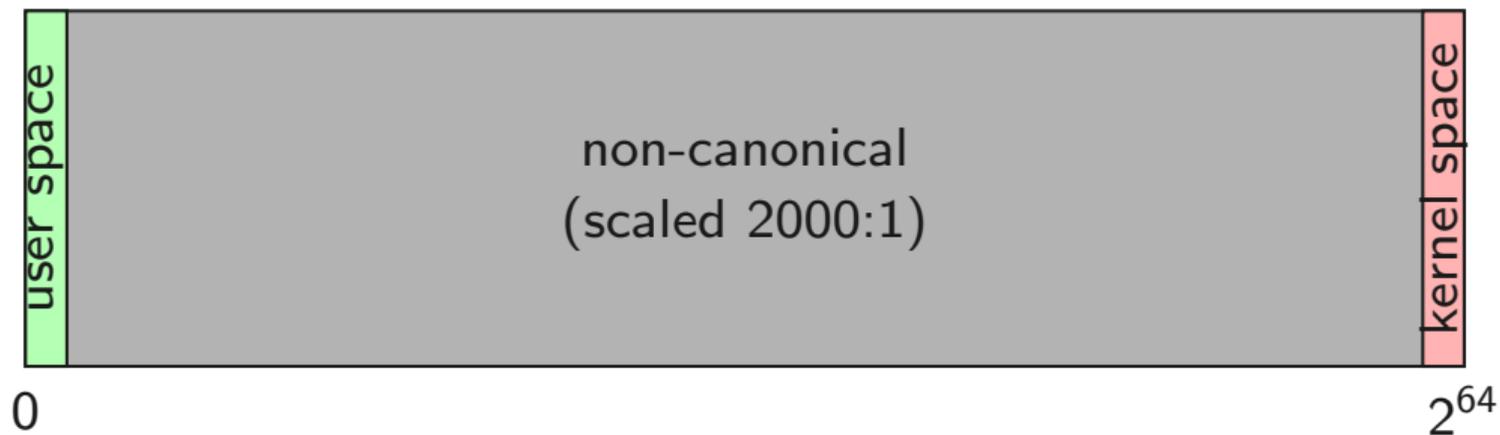
Paging: x86-64 with PML5 and page size 4 KiB



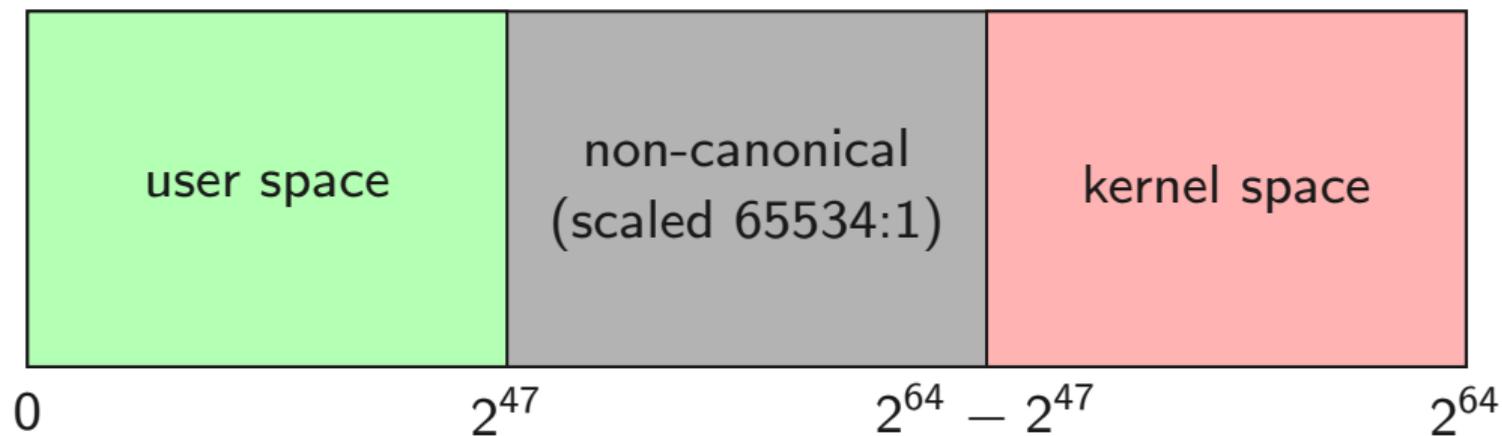
x86-64 Memory Layout (with PML4)



x86-64 Memory Layout (with PML4, scaled)



x86-64 Memory Layout (with PML4, scaled)



Address Translation on x86 processors

- Segmentation and paging

- Segmentation and paging
- 16 K segments, each 4 GB

Intel Pentium

- Segmentation and paging
- 16 K segments, each 4 GB
 - Few segments

Intel Pentium

- Segmentation and paging
- 16 K segments, each 4 GB
 - Few segments
 - Large segments

- Local Descriptor Table LDT

- Local Descriptor Table LDT
 - for each process

- Local Descriptor Table LDT
 - for each process
 - local segments (Code, Data, Stack)

- Local Descriptor Table LDT
 - for each process
 - local segments (Code, Data, Stack)
- Global Descriptor Table GDT

- Local Descriptor Table LDT
 - for each process
 - local segments (Code, Data, Stack)
- Global Descriptor Table GDT
 - for system segments

- Local Descriptor Table LDT
 - for each process
 - local segments (Code, Data, Stack)
- Global Descriptor Table GDT
 - for system segments
 - also for kernel

Segment Registers

- 6 segment registers

Segment Registers

- 6 segment registers
 - CS: Selector for Code Segment

Segment Registers

- 6 segment registers
 - CS: Selector for Code Segment
 - DS: Selector for Data Segment

Segment Registers

- 6 segment registers
 - CS: Selector for Code Segment
 - DS: Selector for Data Segment
 - ES: Selector for Data Segment

Segment Registers

- 6 segment registers
 - CS: Selector for Code Segment
 - DS: Selector for Data Segment
 - ES: Selector for Data Segment
 - FS: Selector for Data Segment

Segment Registers

- 6 segment registers
 - CS: Selector for Code Segment
 - DS: Selector for Data Segment
 - ES: Selector for Data Segment
 - FS: Selector for Data Segment
 - GS: Selector for Data Segment

Segment Registers

- 6 segment registers
 - CS: Selector for Code Segment
 - DS: Selector for Data Segment
 - ES: Selector for Data Segment
 - FS: Selector for Data Segment
 - GS: Selector for Data Segment
 - SS: Selector for Stack Segment

Segment Selector

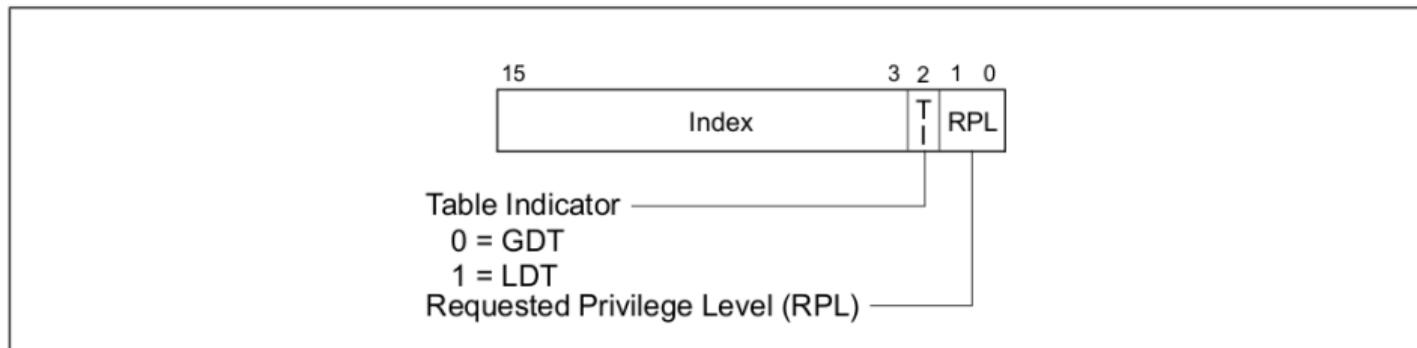


Figure 3-6. Segment Selector

Segment Selector

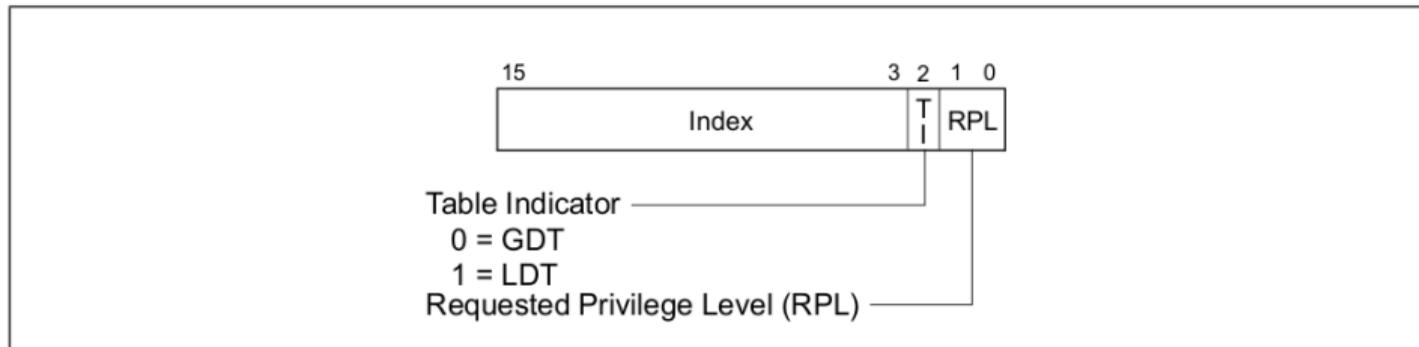


Figure 3-6. Segment Selector

- Null Segment at index 0 → cannot be used

Segment Selector

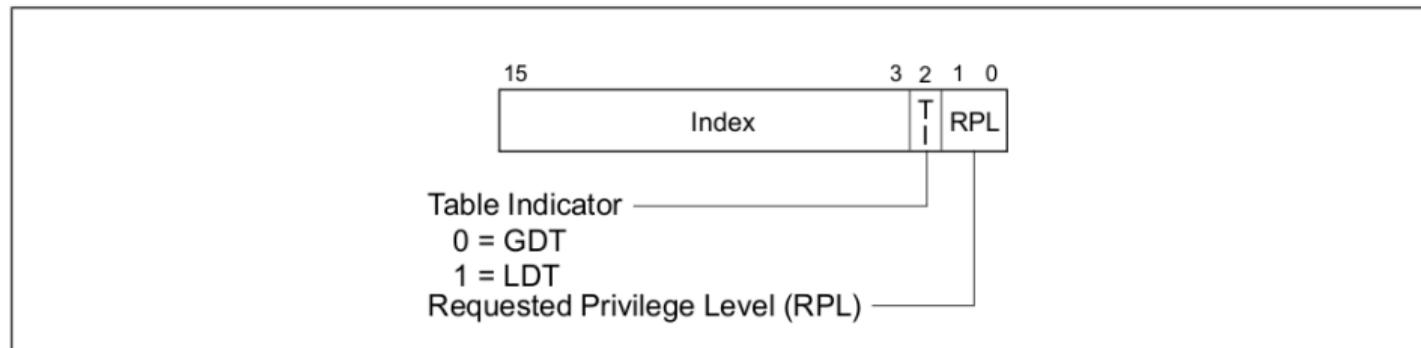


Figure 3-6. Segment Selector

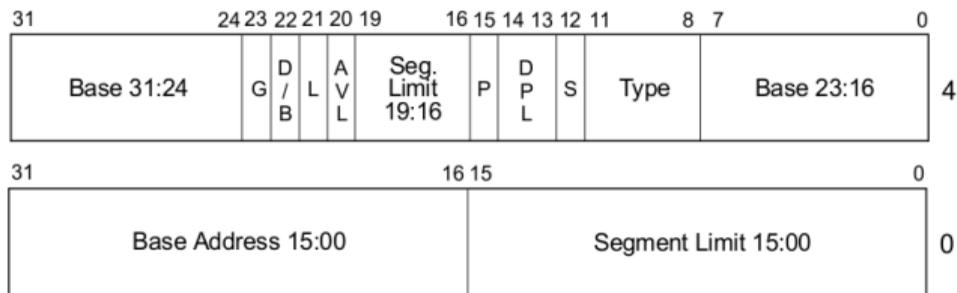
- Null Segment at index 0 → cannot be used
- Modifying a segment register loads corresponding descriptor into an internal CPU register

Hidden Part of Segment Registers

Visible Part		Hidden Part	
Segment Selector	Base Address, Limit, Access Information		
			CS
			SS
			DS
			ES
			FS
			GS

Figure 3-7. Segment Registers

Segment Descriptor



- L — 64-bit code segment (IA-32e mode only)
- AVL — Available for use by system software
- BASE — Segment base address
- D/B — Default operation size (0 = 16-bit segment; 1 = 32-bit segment)
- DPL — Descriptor privilege level
- G — Granularity
- LIMIT — Segment Limit
- P — Segment present
- S — Descriptor type (0 = system; 1 = code or data)
- TYPE — Segment type

Figure 3-8. Segment Descriptor

Address translation

- we start with (selector, offset)

Address translation

- we start with (selector, offset)
- CPU looks for correct descriptor in internal registers

Address translation

- we start with (selector, offset)
- CPU looks for correct descriptor in internal registers
- selector 0 or segment swapped out: interrupt

Address translation

- we start with (selector, offset)
- CPU looks for correct descriptor in internal registers
- selector 0 or segment swapped out: interrupt
- offset exceeds segment size: interrupt

Address translation

- we start with (selector, offset)
- CPU looks for correct descriptor in internal registers
- selector 0 or segment swapped out: interrupt
- offset exceeds segment size: interrupt
- add base field to offset

Address translation

- we start with (selector, offset)
- CPU looks for correct descriptor in internal registers
- selector 0 or segment swapped out: interrupt
- offset exceeds segment size: interrupt
- add base field to offset
 - check limits of course

Address translation

- we start with (selector, offset)
- CPU looks for correct descriptor in internal registers
- selector 0 or segment swapped out: interrupt
- offset exceeds segment size: interrupt
- add base field to offset
 - check limits of course
- result: linear address

Address translation

- we start with (selector, offset)
- CPU looks for correct descriptor in internal registers
- selector 0 or segment swapped out: interrupt
- offset exceeds segment size: interrupt
- add base field to offset
 - check limits of course
- result: linear address
- paging turned off: linear address is physical address

Address translation

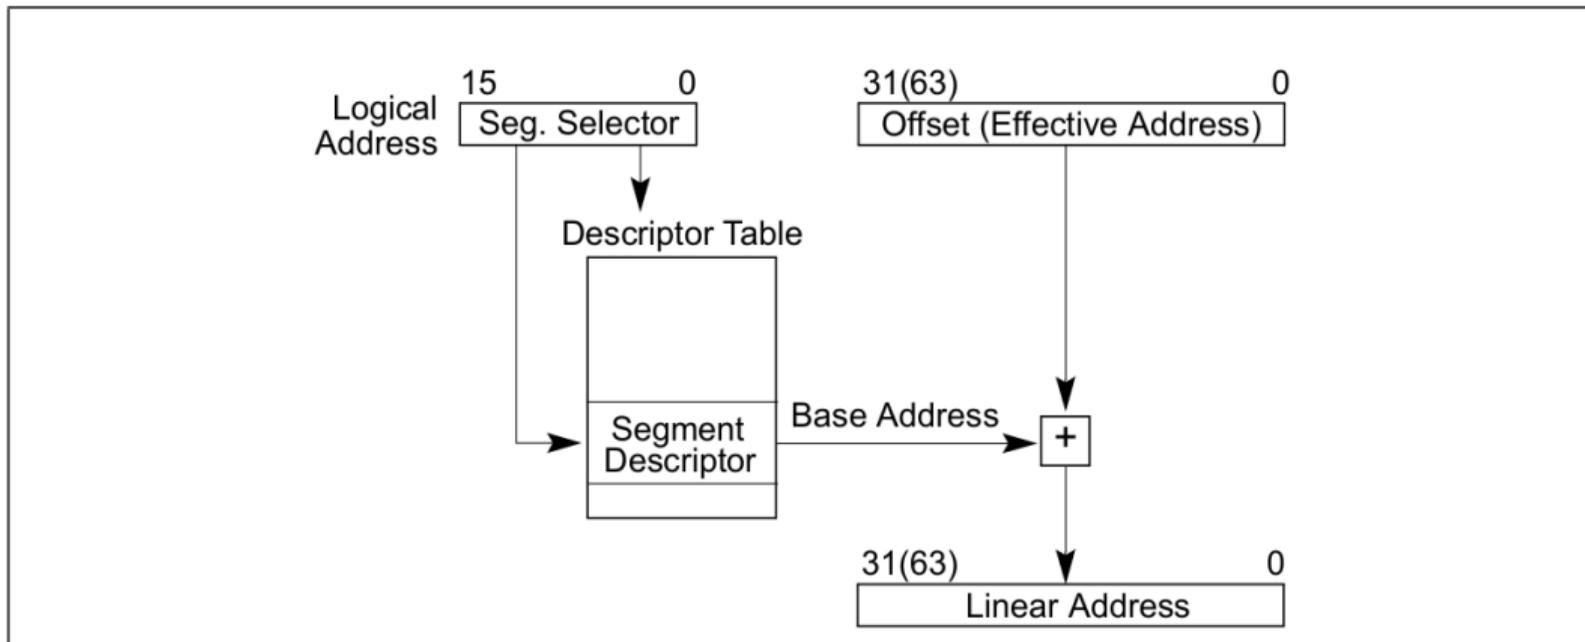


Figure 3-5. Logical Address to Linear Address Translation

Combining Segments and Paging

OSes today have only a very small number of segments:

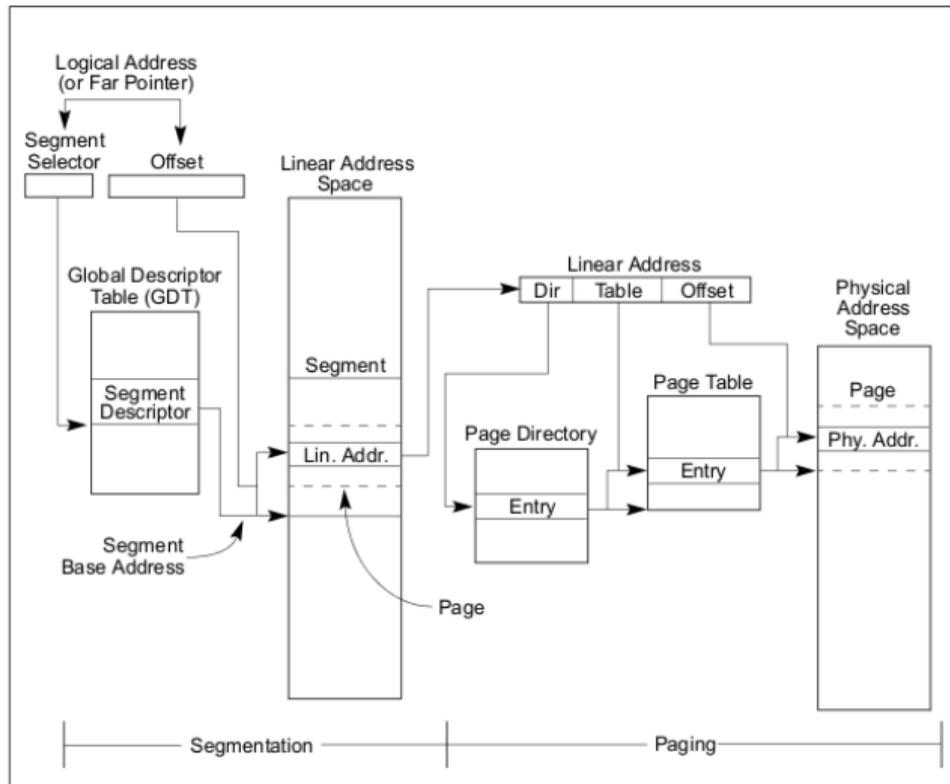


Figure 3-1. Segmentation and Paging

Combining Segments and Paging

OSes today have only a very small number of segments:

- 1 for user code

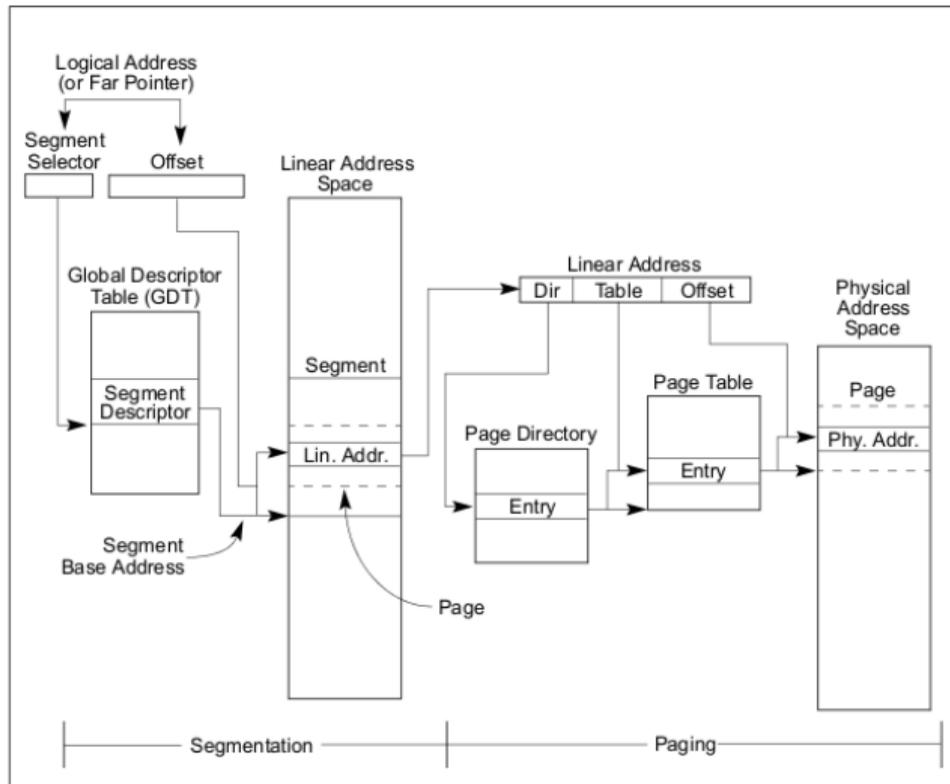


Figure 3-1. Segmentation and Paging

Combining Segments and Paging

OSes today have only a very small number of segments:

- 1 for user code
- 1 for user data

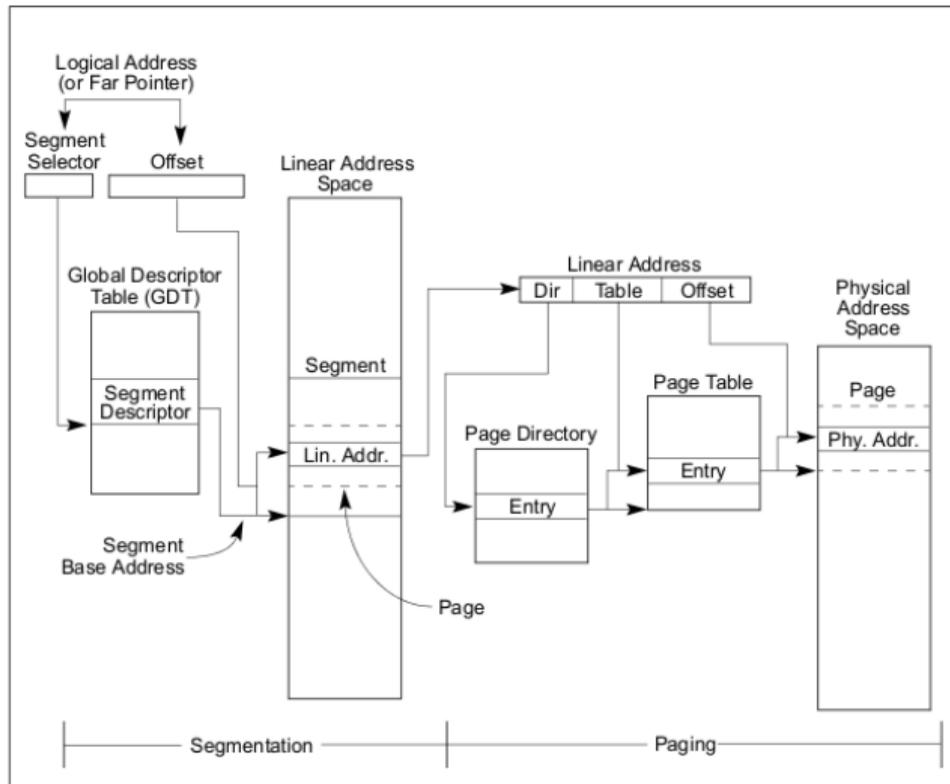


Figure 3-1. Segmentation and Paging

Combining Segments and Paging

OSes today have only a very small number of segments:

- 1 for user code
- 1 for user data
- 1 for user thread local storage

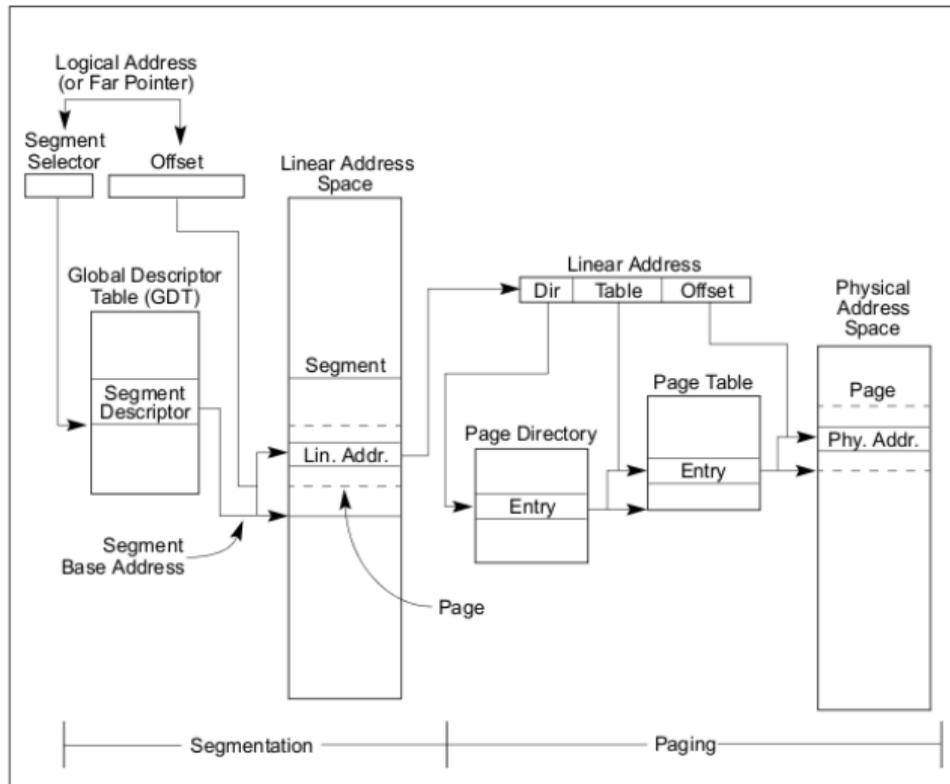


Figure 3-1. Segmentation and Paging

Combining Segments and Paging

OSes today have only a very small number of segments:

- 1 for user code
- 1 for user data
- 1 for user thread local storage
- 1 for kernel code

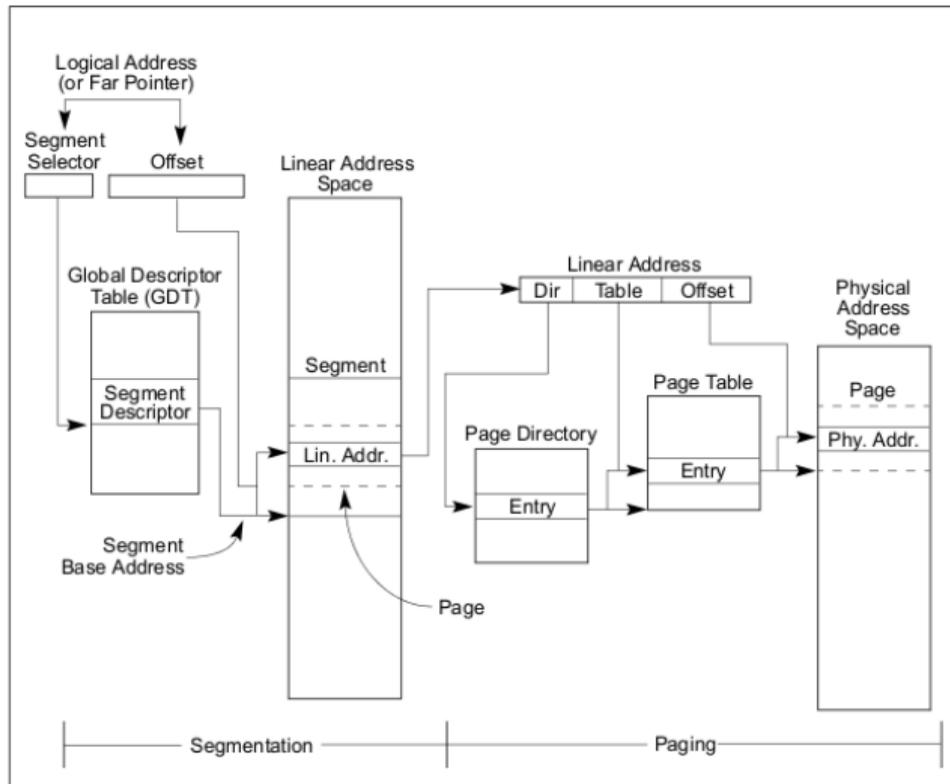


Figure 3-1. Segmentation and Paging

Combining Segments and Paging

OSes today have only a very small number of segments:

- 1 for user code
- 1 for user data
- 1 for user thread local storage
- 1 for kernel code
- 1 for kernel data

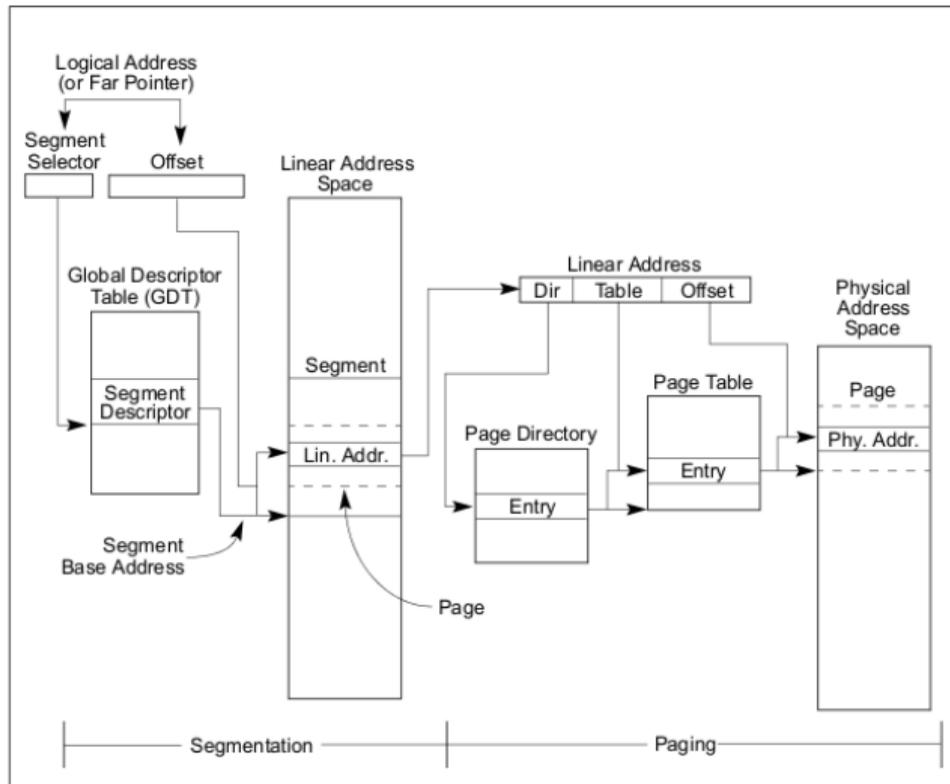


Figure 3-1. Segmentation and Paging

Combining Segments and Paging

OSes today have only a very small number of segments:

- 1 for user code
- 1 for user data
- 1 for user thread local storage
- 1 for kernel code
- 1 for kernel data
- 1 for kernel core local storage

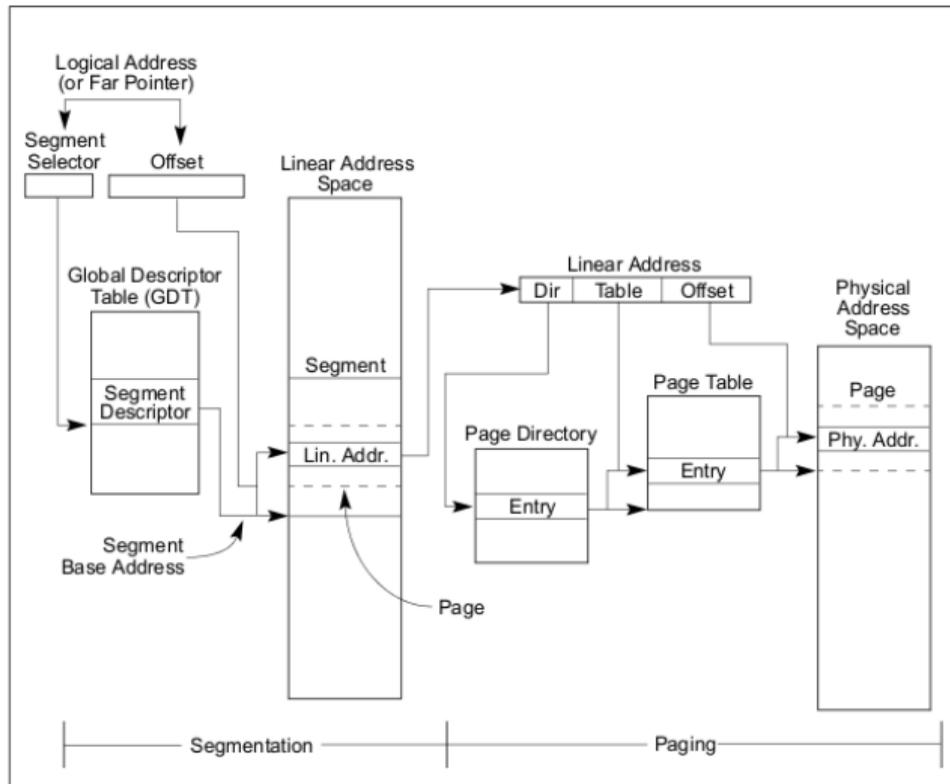


Figure 3-1. Segmentation and Paging

Segments Today

- x86-64 requires segment base to be 0 and limit to be unlimited

Segments Today

- x86-64 requires segment base to be 0 and limit to be unlimited
- not even used anymore to separate code and data

Segments Today

- x86-64 requires segment base to be 0 and limit to be unlimited
- not even used anymore to separate code and data
- most OSes today only use segments to determine the privilege level

Take Aways

Virtual memory

Take Aways

Virtual memory

- is based on Segmentation and Paging

Take Aways

Virtual memory

- is based on Segmentation and Paging
- enables effective protection mechanisms

Take Aways

Virtual memory

- is based on Segmentation and Paging
- enables effective protection mechanisms
-

Take Aways

Virtual memory

- is based on Segmentation and Paging
- enables effective protection mechanisms
-
- enables sparse address spaces

