

*Operating
Systems:
Internals
and Design
Principles*

Chapter 8 Virtual Memory

Eighth Edition
William Stallings

Virtual memory	A storage allocation scheme in which secondary memory can be addressed as though it were part of main memory. The addresses a program may use to reference memory are distinguished from the addresses the memory system uses to identify physical storage sites, and program-generated addresses are translated automatically to the corresponding machine addresses. The size of virtual storage is limited by the addressing scheme of the computer system and by the amount of secondary memory available and not by the actual number of main storage locations.
Virtual address	The address assigned to a location in virtual memory to allow that location to be accessed as though it were part of main memory.
Virtual address space	The virtual storage assigned to a process.
Address space	The range of memory addresses available to a process.
Real address	The address of a storage location in main memory.

Table 8.1 Virtual Memory Terminology

Hardware and Control Structures

- Two characteristics fundamental to memory management:
 - 1) all memory references are logical addresses that are dynamically translated into physical addresses at run time
 - 2) a process may be broken up into a number of pieces that don't need to be contiguously located in main memory during execution
- If these two characteristics are present, it is not necessary that all of the pages or segments of a process be in main memory during execution

Execution of a Process

- Operating system brings into main memory a few pieces of the program
- Resident set
 - portion of process that is in main memory
- An interrupt is generated when an address is needed that is not in main memory
- Operating system places the process in a blocking state



Continued . . .

Execution of a Process

- Piece of process that contains the logical address is brought into main memory
 - operating system issues a disk I/O Read request
 - another process is dispatched to run while the disk I/O takes place
 - an interrupt is issued when disk I/O is complete, which causes the operating system to place the affected process in the Ready state



Implications


- More processes may be maintained in main memory
 - only load in some of the pieces of each process
 - with so many processes in main memory, it is very likely a process will be in the Ready state at any particular time
- A process may be larger than all of main memory



Real and Virtual Memory

Real
memory

main memory,
the actual
RAM



Virtual
memory

memory on disk


allows for effective
multiprogramming
and relieves the user
of tight constraints
of main memory

Simple Paging	Virtual Memory Paging	Simple Segmentation	Virtual Memory Segmentation
Main memory partitioned into small fixed-size chunks called frames		Main memory not partitioned	
Program broken into pages by the compiler or memory management system		Program segments specified by the programmer to the compiler (i.e., the decision is made by the programmer)	
Internal fragmentation within frames		No internal fragmentation	
No external fragmentation		External fragmentation	
Operating system must maintain a page table for each process showing which frame each page occupies		Operating system must maintain a segment table for each process showing the load address and length of each segment	
Operating system must maintain a free frame list		Operating system must maintain a list of free holes in main memory	
Processor uses page number, offset to calculate absolute address		Processor uses segment number, offset to calculate absolute address	
All the pages of a process must be in main memory for process to run, unless overlays are used	Not all pages of a process need be in main memory frames for the process to run. Pages may be read in as needed	All the segments of a process must be in main memory for process to run, unless overlays are used	Not all segments of a process need be in main memory for the process to run. Segments may be read in as needed
	Reading a page into main memory may require writing a page out to disk		Reading a segment into main memory may require writing one or more segments out to disk

Table 8.2

Characteristics of Paging and Segmentation

Thrashing



A state in which the system spends most of its time swapping process pieces rather than executing instructions

The diagram features a large, light purple circle on the left side. Inside this circle is a smaller, solid dark red circle containing text. To the right of this inner circle is another solid dark red circle, also containing text. The two dark red circles overlap slightly. The background of the slide is a light beige, textured surface.

To avoid this, the operating system tries to guess, based on recent history, which pieces are least likely to be used in the near future

Principle of Locality

- Program and data references within a process tend to cluster
- Only a few pieces of a process will be needed over a short period of time
- Therefore it is possible to make intelligent guesses about which pieces will be needed in the future
- Avoids thrashing



Support Needed for Virtual Memory

For virtual memory to be practical and effective:


- hardware must support paging and segmentation
- operating system must include software for managing the movement of pages and/or segments between secondary memory and main memory

Paging

- The term *virtual memory* is usually associated with systems that employ paging
- Use of paging to achieve virtual memory was first reported for the Atlas computer
- Each process has its own page table
 - each page table entry contains the frame number of the corresponding page in main memory



Virtual Address



Page Table Entry



(a) Paging only

Virtual Address




Segment Table Entry



(b) Segmentation only

Virtual Address



Segment Table Entry



Page Table Entry



P= present bit
M = Modified bit

(c) Combined segmentation and paging

Figure 8.1 Typical Memory Management Formats






Figure 8.2 Address Translation in a Paging System

4-kbyte root
page table



4-M byte user
page table



4-Gbyte user
address space



Figure 8.3 A Two-Level Hierarchical Page Table




Figure 8.4 Address Translation in a Two-Level Paging System


Inverted Page Table

- Page number portion of a virtual address is mapped into a hash value
 - hash value points to inverted page table
- Fixed proportion of real memory is required for the tables regardless of the number of processes or virtual pages supported
- Structure is called *inverted* because it indexes page table entries by frame number rather than by virtual page number



Inverted Page Table

Each entry in the page table includes:



Translation Lookaside Buffer (TLB)

- Each virtual memory reference can cause two physical memory accesses:
 - one to fetch the page table entry
 - one to fetch the data
- To overcome the effect of doubling the memory access time, most virtual memory schemes make use of a special high-speed cache called a *translation lookaside buffer*




Figure 8.6 Use of a Translation Lookaside Buffer





Figure 8.7 Operation of Paging and Translation Lookaside Buffer (TLB) [FURH87]

Associative Mapping


- The TLB only contains some of the page table entries so we cannot simply index into the TLB based on page number
 - each TLB entry must include the page number as well as the complete page table entry
- The processor is equipped with hardware that allows it to interrogate simultaneously a number of TLB entries to determine if there is a match on page number



Virtual Address

Page# Offset

5	502
---	-----



Page Table

(a) Direct mapping


37	502
----	-----

Frame# Offset
Real Address

Virtual Address

Page# Offset

5	502
---	-----



Translation Lookaside Buffer

(b) Associative mapping

37	502
----	-----

Frame# Offset
Real Address

Figure 8.8 Direct Versus Associative Lookup for Page Table Entries

TLB Operation






Figure 8.9 Translation Lookaside Buffer and Cache Operation

Page Size

- The smaller the page size, the lesser the amount of internal fragmentation
 - however, more pages are required per process
 - more pages per process means larger page tables
 - for large programs in a heavily multiprogrammed environment some portion of the page tables of active processes must be in virtual memory instead of main memory
 - the physical characteristics of most secondary-memory devices favor a larger page size for more efficient block transfer of data



(a) Page Size



(b) Number of Page Frames Allocated

P = size of entire process
 W = working set size
 N = total number of pages in process

Figure 8.10 Typical Paging Behavior of a Program

Computer	Page Size
Atlas	512 48-bit words
Honeywell-Multics	1024 36-bit words
IBM 370/XA and 370/ESA	4 Kbytes
VAX family	512 bytes
IBM AS/400	512 bytes
DEC Alpha	8 Kbytes
MIPS	4 Kbytes to 16 Mbytes
UltraSPARC	8 Kbytes to 4 Mbytes
Pentium	4 Kbytes or 4 Mbytes
IBM POWER	4 Kbytes
Itanium	4 Kbytes to 256 Mbytes

Table 8.3

**Example
Page
Sizes**

Page Size

the design issue of page size is related to the size of physical main memory and program size



main memory is getting larger and address space used by applications is also growing



most obvious on personal computers where applications are becoming increasingly complex

- Contemporary programming techniques used in large programs tend to decrease the locality of references within a process

Segmentation


- Segmentation allows the programmer to view memory as consisting of multiple address spaces or segments

Advantages:

- simplifies handling of growing data structures
- allows programs to be altered and recompiled independently
- lends itself to sharing data among processes
- lends itself to protection

Segment Organization

- Each segment table entry contains the starting address of the corresponding segment in main memory and the length of the segment
- A bit is needed to determine if the segment is already in main memory
- Another bit is needed to determine if the segment has been modified since it was loaded in main memory







Figure 8.11 Address Translation in a Segmentation System

Combined Paging and Segmentation

In a combined paging/segmentation system a user's address space is broken up into a number of segments. Each segment is broken up into a number of fixed-sized pages which are equal in length to a main memory frame

Segmentation is visible to the programmer


Paging is transparent to the programmer



Virtual Address



Segment Table Entry



Page Table Entry




P= present bit
M = Modified bit

(c) Combined segmentation and paging

Protection and Sharing

- Segmentation lends itself to the implementation of protection and sharing policies
- Each entry has a base address and length so inadvertent memory access can be controlled
- Sharing can be achieved by segments referencing multiple processes






Figure 8.13 Protection Relationships Between Segments

Operating System Software

The design of the memory management portion of an operating system depends on three fundamental areas of choice:

- whether or not to use virtual memory techniques
- the use of paging or segmentation or both
- the algorithms employed for various aspects of memory management

<p>Fetch Policy</p> <ul style="list-style-type: none"> Demand paging Prepaging <p>Placement Policy</p> <p>Replacement Policy</p> <ul style="list-style-type: none"> Basic Algorithms <ul style="list-style-type: none"> Optimal Least recently used (LRU) First-in-first-out (FIFO) Clock Page Buffering 	<p>Resident Set Management</p> <ul style="list-style-type: none"> Resident set size <ul style="list-style-type: none"> Fixed Variable Replacement Scope <ul style="list-style-type: none"> Global Local <p>Cleaning Policy</p> <ul style="list-style-type: none"> Demand Precleaning <p>Load Control</p> <ul style="list-style-type: none"> Degree of multiprogramming
--	--

Table 8.4 Operating System Policies for Virtual Memory

Fetch Policy

- Determines when a page should be brought into memory



Two main types:


Demand
Paging

Prepaging

Demand Paging

■ Demand Paging

- only brings pages into main memory when a reference is made to a location on the page
- many page faults when process is first started
- principle of locality suggests that as more and more pages are brought in, most future references will be to pages that have recently been brought in, and page faults should drop to a very low level



Prepaging

- **Prepaging**
 - pages other than the one demanded by a page fault are brought in
 - exploits the characteristics of most secondary memory devices
 - if pages of a process are stored contiguously in secondary memory it is more efficient to bring in a number of pages at one time
 - ineffective if extra pages are not referenced
 - should not be confused with “swapping”

Placement Policy

- Determines where in real memory a process piece is to reside
- Important design issue in a segmentation system
- Paging or combined paging with segmentation placing is irrelevant because hardware performs functions with equal efficiency
- For NUMA systems an automatic placement strategy is desirable

Replacement Policy

- Deals with the selection of a page in main memory to be replaced when a new page must be brought in
 - objective is that the page that is removed be the page least likely to be referenced in the near future
- The more elaborate the replacement policy the greater the hardware and software overhead to implement it

Frame Locking

- When a frame is locked the page currently stored in that frame may not be replaced
 - kernel of the OS as well as key control structures are held in locked frames
 - I/O buffers and time-critical areas may be locked into main memory frames
 - locking is achieved by associating a lock bit with each frame



Basic Algorithms



Algorithms used for the selection of a page to replace:

- Optimal
- Least recently used (LRU)
- First-in-first-out (FIFO)
- Clock

Page address stream

2 3 2 1 5 2 4 5 3 2 5 2

OPT

2	2	2	2	2	2	4	4	4	2	2	2
	3	3	3	3	3	3	3	3	3	3	3
			1	5	5	5	5	5	5	5	5
				F		F			F		

LRU

2	2	2	2	2	2	2	2	3	3	3	3
	3	3	3	5	5	5	5	5	5	5	5
			1	1	1	4	4	4	2	2	2
				F		F		F	F		

FIFO

2	2	2	2	5	5	5	5	3	3	3	3
	3	3	3	3	2	2	2	2	2	5	5
			1	1	1	4	4	4	4	4	2
				F	F	F		F		F	F

CLOCK

2*	2*	2*	2*	5*	5*	5*	5*	3*	3*	3*	3*
	3*	3*	3*	3	2*	2*	2*	2	2*	2	2*
			1*	1	1	4*	4*	4	4	5*	5*
				F	F	F		F		F	

F = page fault occurring after the frame allocation is initially filled

Figure 8.14 Behavior of Four Page-Replacement Algorithms


Least Recently Used (LRU)

- Replaces the page that has not been referenced for the longest time
- By the principle of locality, this should be the page least likely to be referenced in the near future
- Difficult to implement
 - one approach is to tag each page with the time of last reference
 - this requires a great deal of overhead




First-in-First-out (FIFO)


- Treats page frames allocated to a process as a circular buffer
- Pages are removed in round-robin style
 - simple replacement policy to implement
- Page that has been in memory the longest is replaced




Clock Policy

- Requires the association of an additional bit with each frame
 - referred to as the *use* bit
- When a page is first loaded in memory or referenced, the use bit is set to 1
- The set of frames is considered to be a circular buffer
- Any frame with a use bit of 1 is passed over by the algorithm
- Page frames visualized as laid out in a circle





(a) State of buffer just prior to a page replacement



(b) State of buffer just after the next page replacement

Figure 8.15 Example of Clock Policy Operation





Figure 8.16 Comparison of Fixed-Allocation, Local Page Replacement Algorithms


Page Buffering

- Improves paging performance and allows the use of a simpler page replacement policy



Replacement Policy and Cache Size

- With large caches, replacement of pages can have a performance impact
 - if the page frame selected for replacement is in the cache, that cache block is lost as well as the page that it holds
 - in systems using page buffering, cache performance can be improved with a policy for page placement in the page buffer
 - most operating systems place pages by selecting an arbitrary page frame from the page buffer



Resident Set Management

- The OS must decide how many pages to bring into main memory
 - the smaller the amount of memory allocated to each process, the more processes can reside in memory
 - small number of pages loaded increases page faults
 - beyond a certain size, further allocations of pages will not effect the page fault rate



Resident Set Size

Fixed-allocation

- gives a process a fixed number of frames in main memory within which to execute
 - when a page fault occurs, one of the pages of that process must be replaced

Variable-allocation

- allows the number of page frames allocated to a process to be varied over the lifetime of the process

Replacement Scope

- The scope of a replacement strategy can be categorized as *global* or *local*
 - both types are activated by a page fault when there are no free page frames

Local

- chooses only among the resident pages of the process that generated the page fault

Global

- considers all unlocked pages in main memory

	Local Replacement	Global Replacement
Fixed Allocation	<ul style="list-style-type: none"> •Number of frames allocated to a process is fixed. •Page to be replaced is chosen from among the frames allocated to that process. 	<ul style="list-style-type: none"> •Not possible.
Variable Allocation	<ul style="list-style-type: none"> •The number of frames allocated to a process may be changed from time to time to maintain the working set of the process. •Page to be replaced is chosen from among the frames allocated to that process. 	<ul style="list-style-type: none"> •Page to be replaced is chosen from all available frames in main memory; this causes the size of the resident set of processes to vary.

Table 8.5 Resident Set Management

Fixed Allocation, Local Scope

- Necessary to decide ahead of time the amount of allocation to give a process
- If allocation is too small, there will be a high page fault rate

If allocation is too large, there will be too few programs in main memory

- increased processor idle time
- increased time spent in swapping

Variable Allocation

Global Scope

- Easiest to implement
 - adopted in a number of operating systems
- OS maintains a list of free frames
- Free frame is added to resident set of process when a page fault occurs
- If no frames are available the OS must choose a page currently in memory
- One way to counter potential problems is to use page buffering

Variable Allocation

Local Scope

- When a new process is loaded into main memory, allocate to it a certain number of page frames as its resident set
- When a page fault occurs, select the page to replace from among the resident set of the process that suffers the fault
- Reevaluate the allocation provided to the process and increase or decrease it to improve overall performance



Variable Allocation

Local Scope

- Decision to increase or decrease a resident set size is based on the assessment of the likely future demands of active processes

Key elements:

- criteria used to determine resident set size
- the timing of changes

Sequence of Page References	Window Size, D			
	2	3	4	5
24	24	24	24	24
15	24 15	24 15	24 15	24 15
18	15 18	24 15 18	24 15 18	24 15 18
23	18 23	15 18 23	24 15 18 23	24 15 18 23
24	23 24	18 23 24	•	•
17	24 17	23 24 17	18 23 24 17	15 18 23 24 17
18	17 18	24 17 18	•	18 23 24 17
24	18 24	•	24 17 18	•
18	•	18 24	•	24 17 18
17	18 17	24 18 17	•	•
17	17	18 17	•	•
15	17 15	17 15	18 17 15	24 18 17 15
24	15 24	17 15 24	17 15 24	•
17	24 17	•	•	17 15 24
24	•	24 17	•	•
18	24 18	17 24 18	17 24 18	15 17 24 18

Figure 8.17 Working Set of Process as Defined by Window Size




Figure 8.18 Typical Graph of Working Set Size [MAEK87]

Page Fault Frequency (PFF)

- Requires a use bit to be associated with each page in memory
- Bit is set to 1 when that page is accessed
- When a page fault occurs, the OS notes the virtual time since the last page fault for that process
- Does not perform well during the transient periods when there is a shift to a new locality

Variable-Interval Sampled Working Set (VSWS)

- Evaluates the working set of a process at sampling instances based on elapsed virtual time
- Driven by three parameters:

the minimum
duration of the
sampling
interval

the maximum
duration of the
sampling
interval

the number of
page faults that
are allowed to
occur between
sampling
instances

Cleaning Policy

- Concerned with determining when a modified page should be written out to secondary memory

Demand Cleaning

a page is written out to secondary memory only when it has been selected for replacement




Precleaning

allows the writing of pages in batches

Load Control

- Determines the number of processes that will be resident in main memory
 - *multiprogramming* level
- Critical in effective memory management
- Too few processes, many occasions when all processes will be blocked and much time will be spent in swapping
- Too many processes will lead to thrashing






Figure 8.19 Multiprogramming Effects

Process Suspension

- If the degree of multiprogramming is to be reduced, one or more of the currently resident processes must be swapped out

Six possibilities exist:

- lowest-priority process
- faulting process
- last process activated
- process with the smallest resident set
- largest process
- process with the largest remaining execution window

UNIX

- Intended to be machine independent so its memory management schemes will vary
 - early UNIX: variable partitioning with no virtual memory scheme
 - current implementations of UNIX and Solaris make use of

**SVR4 and Solaris use
two separate schemes:**

- paging system
- kernel memory allocator

Paging System and Kernel Memory Allocator

Paging System

provides a virtual memory capability that allocates page frames in main memory to processes

allocates page frames to disk block buffers

Kernel Memory Allocator

allocates memory for the kernel

Page frame number	Age	Copy on write	Modify	Reference	Valid	Protect
-------------------	-----	---------------	--------	-----------	-------	---------

(a) Page table entry

Swap device number	Device block number	Type of storage
--------------------	---------------------	-----------------

(b) Disk block descriptor

Page state	Reference count	Logical device	Block number	Pfdata pointer
------------	-----------------	----------------	--------------	----------------

(c) Page frame data table entry

Reference count	Page/storage unit number
-----------------	--------------------------

(d) Swap-use table entry

Figure 8.20 UNIX SVR4 Memory Management Formats

Page Table Entry

Page frame number

Refers to frame in real memory.

Age

Indicates how long the page has been in memory without being referenced. The length and contents of this field are processor dependent.

Copy on write

Set when more than one process shares a page. If one of the processes writes into the page, a separate copy of the page must first be made for all other processes that share the page. This feature allows the copy operation to be deferred until necessary and avoided in cases where it turns out not to be necessary.

Modify

Indicates page has been modified.

Reference

Indicates page has been referenced. This bit is set to 0 when the page is first loaded and may be periodically reset by the page replacement algorithm.

Valid

Indicates page is in main memory.

Protect

Indicates whether write operation is allowed.

Disk Block Descriptor

Swap device number

Logical device number of the secondary device that holds the corresponding page. This allows more than one device to be used for swapping.

Device block number

Block location of page on swap device.

Type of storage

Storage may be swap unit or executable file. In the latter case, there is an indication as to whether the virtual memory to be allocated should be cleared first.

Table 8.6

UNIX SVR4 Memory Management Parameters (page 1 of 2)

(Table can be found on page 379
in the textbook)

Page Frame Data Table Entry

Page state

Indicates whether this frame is available or has an associated page. In the latter case, the status of the page is specified: on swap device, in executable file, or DMA in progress.

Reference count

Number of processes that reference the page.

Logical device

Logical device that contains a copy of the page.

Block number

Block location of the page copy on the logical device.

Pfdata pointer

Pointer to other pfdata table entries on a list of free pages and on a hash queue of pages.

Swap-Use Table Entry

Reference count

Number of page table entries that point to a page on the swap device.

Page/storage unit number

Page identifier on storage unit.

Table 8.6

UNIX SVR4 Memory Management Parameters (page 2 of 2)

(Table can be found on page 379 in the textbook)

Page Replacement

- The page frame data table is used for page replacement
- Pointers are used to create lists within the table
 - all available frames are linked together in a list of free frames available for bringing in pages
 - when the number of available frames drops below a certain threshold, the kernel will steal a number of frames to compensate






Figure 8.21 Two-Handed Clock Page-Replacement Algorithm


Kernel Memory Allocator

- The kernel generates and destroys small tables and buffers frequently during the course of execution, each of which requires dynamic memory allocation.
- Most of these blocks are significantly smaller than typical pages (therefore paging would be inefficient)
- Allocations and free operations must be made as fast as possible



Lazy Buddy

- Technique adopted for SVR4
- UNIX often exhibits steady-state behavior in kernel memory demand
 - i.e. the amount of demand for blocks of a particular size varies slowly in time
- Defers coalescing until it seems likely that it is needed, and then coalesces as many blocks as possible



Initial value of D_j is 0

After an operation, the value of D_j is updated as follows

(I) if the next operation is a block allocate request:

if there is any free block, select one to allocate

if the selected block is locally free

then $D_j := D_j + 2$

else $D_j := D_j + 1$

otherwise

first get two blocks by splitting a larger one into two (recursive operation)

allocate one and mark the other locally free

D_j remains unchanged (but D may change for other block sizes because of the recursive call)

(II) if the next operation is a block free request

Case $D_j \geq 2$

mark it locally free and free it locally

$D_j := D_j - 2$

Case $D_j = 1$

mark it globally free and free it globally; coalesce if possible

$D_j := 0$

Case $D_j = 0$

mark it globally free and free it globally; coalesce if possible

select one locally free block of size 2^i and free it globally; coalesce if possible

$D_j := 0$

Figure 8.22 Lazy Buddy System Algorithm

Linux

Memory Management


- Shares many characteristics with UNIX
- Is quite complex

Two main
aspects

- process virtual memory
- kernel memory allocation

Linux Virtual Memory

■ Three level page table structure:






Figure 8.23 Address Translation in Linux Virtual Memory Scheme

Linux Page Replacement

- Based on the clock algorithm
- The use bit is replaced with an 8-bit age variable
 - incremented each time the page is accessed
- Periodically decrements the age bits
 - a page with an age of 0 is an “old” page that has not been referenced in some time and is the best candidate for replacement
- A form of least frequently used policy




Figure 8.24 Linux Page Reclaiming

Kernel Memory Allocation

- Kernel memory capability manages physical main memory page frames
 - primary function is to allocate and deallocate frames for particular uses



Possible owners of a frame include:

- user-space processes
- dynamically allocated kernel data
- static kernel code
- page cache

- A buddy algorithm is used so that memory for the kernel can be allocated and deallocated in units of one or more pages
- Page allocator alone would be inefficient because the kernel requires small short-term memory chunks in odd sizes
- Slab allocation
 - used by Linux to accommodate small chunks

Windows Memory Management

- Virtual memory manager controls how memory is allocated and how paging is performed
- Designed to operate over a variety of platforms
- Uses page sizes ranging from 4 Kbytes to 64 Kbytes



Windows Virtual Address Map

- On 32 bit platforms each user process sees a separate 32 bit address space allowing 4 Gbytes of virtual memory per process
 - by default half is reserved for the OS
- Large memory intensive applications run more effectively using 64-bit Windows
- Most modern PCs use the AMD64 processor architecture which is capable of running as either a 32-bit or 64-bit system






Figure 8.25 Windows Default 32-bit Virtual Address Space


Windows Paging

- On creation, a process can make use of the entire user space of almost 2 Gbytes
- This space is divided into fixed-size pages managed in contiguous regions allocated on 64 Kbyte boundaries
- Regions may be in one of three states:



Resident Set Management System

- Windows uses variable allocation, local scope
- When activated, a process is assigned a data structure to manage its working set
- Working sets of active processes are adjusted depending on the availability of main memory



Android Memory Management

- Android includes a number of extensions to the normal Linux kernel memory management facility
- These include:
 - ASHMem
 - this feature provides anonymous shared memory, which abstracts memory as file descriptors
 - a file descriptor can be passed to another process to share memory
 - Pmem
 - this feature allocates virtual memory so that it is physically contiguous
 - this feature is useful for hardware that does not support virtual memory
 - Low Memory Killer
 - this feature enables the system to notify an app or apps that they need to free up memory
 - if an app does not cooperate, it is terminated

Summary

- Hardware and control structures
 - Locality and virtual memory
 - Paging
 - Segmentation
 - Combined paging and segmentation
 - Protection and sharing
- OS software
 - Fetch policy
 - Placement policy
 - Replacement policy
 - Resident set management
 - Cleaning policy
 - Load control
- UNIX and Solaris memory management
 - Paging system
 - Kernel memory allocator
- Linux memory management
 - Linux virtual memory
 - Kernel memory allocation
- Windows memory management
 - Windows virtual address map
 - Windows paging
 - Windows 8 swapping
- Android memory management